

Table of Contents

Editors' Introduction IRENE LO and SAM TAGGART	1
Statements from SIGecom Officers Candidates IRENE LO and SAM TAGGART	3
SIGecom Winter Meeting 2023 Highlights EMILY RYU, CHENGHAN ZHOU and MARYAM BAHRANI	5
Report on EAAMO 2022 ELENA FALCETTONI, DINA MACHUVE, BRYAN WILDER and ANGELA ZHOU	14
Recent Developments in Pandora's Box Problem: Variants and Applications HEDYEH BEYHAGHI and LINDA CAI	20
Tractable Choice MODIBO K. CAMARA	35
A Proof of the Nisan-Ronen Conjecture — An Overview GEORGE CHRISTODOULOU, ELIAS KOUTSOUPIS and ANNAMARIA KOVACS	42
Deep Reinforcement Learning for Economics: Progress and Challenges ETAN A. GREEN and E. BARRY PLUNKETT	49
Mechanism Design with Predictions: An Annotated Reading List ERIC BALKANSKI, VASILIS GKATZELIS and XIZHI TAN	54
Solution to Exchanges 20.1 Puzzle: Communicating to Plan Noam Nisan's 60th Birthday Workshop ALEC SUN	58
Puzzle: Does Occasional Simulation Enable Cooperation? (Puzzle in honor of Joe Halpern's 70th birthday) VINCENT CONITZER	62

ACM SIGecom Exchanges, Vol. 21, No. 1, June 2023

Editors-in-Chief: Irene Lo and Sam Taggart

Information Director: Yannai A. Gonczarowski

ACM Staff: Irene Frawley

Notice to Contributing Authors to SIG Newsletters

As a contributing author, you retain copyright to your article. ACM will refer all requests for republication directly to you.

By submitting your article for distribution in any newsletter of the ACM Special Interest Groups, you hereby grant to ACM the following non-exclusive, perpetual, worldwide rights:

- to publish your work online or in print on condition of acceptance by the editor
- to include the article in the ACM Digital Library and in any Digital Library-related services
- to allow users to make a personal copy of the article for noncommercial, educational, or research purposes
- to upload your video and other supplemental material to the ACM Digital Library, the ACM YouTube channel, and the SIG newsletter site

Furthermore, you affirm that:

- if third-party materials were used in your published work, supplemental material, or video, that you have the necessary permissions to use those third-party materials in your work

Editors' Introduction

IRENE LO

Stanford University

and

SAM TAGGART

Oberlin College

This summer issue of SIGecom Exchanges has an exciting mix of news updates and technical content. The issue starts with statements from this year's slate of candidates for SIGecom officer positions. It has two event summaries, for the 2023 SIGecomm Winter Meeting and the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO'22). There is a survey, three research letters, and an annotated reading list. Finally, puzzles: a solution to the puzzle published in the Summer 2022 issue of the Exchanges in honor of Noam Nisan's 60th birthday, and a new puzzle in honor of Joe Halpern's 70th birthday.

SIGecom sponsors several events throughout the year. This issue highlights two. The third SIGecom Winter Meeting took place virtually in February 2023 on the topic of Web3/Blockchain/Cryptocurrencies. Graduate students Emily Ryu, Chenghan Zhou, and Maryam Bahrani provide an excellent overview of the event. In their coverage, they summarize an introductory presentation and panel on *what* blockchain and Web3 are and *why* researchers should work on problems in this space. They also provide excerpts from a fireside chat with Tim Roughgarden, an overview of a social activity where participants could mint their own NFT, and a summary of an NFT discussion, treating the Bored Ape Yacht Club as a case study.

The second ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO'22) took place in October 2022 at George Mason University in Arlington, Virginia. We invited the program chairs Elena Falcettoni, Dina Machuve, Bryan Wilder, and Angela Zhou to contribute a report summarizing the technical program, events, and organization of the conference. The report summarizes emergent themes in this year's program, and highlights the work that won awards. It should be of interest to researchers interested in equity, access, and social issues across a wide range of disciplines.

Hedyeh Beyhaghi and Linda Cai authored a comprehensive survey on Pandora's box problem for sequential search with costly inspections. Their survey first presents the canonical version of the problem by Weitzman in 1979. They then overview the wide range of recent extensions to the model, which impose additional combinatorial structure on the problem, constrain the searcher's information, or limit their adaptivity, among other things. The survey also covers applications of the model, including to mechanism design and matching markets.

Modibo Camara, winner of the 2022 EC Best Paper with a Student Lead Author Award, 2022 EC Exemplary Theory Track Paper Award, and 2023 ACM SIGecom

Author's address: ilo@stanford.edu, italgam@cs.technion.ac.il.

Doctoral Dissertation Award, provides an overview of his groundbreaking theory of *tractable choice*. Using his EC'22 paper as a case study, he argues that economic assumptions of rational choice and a computational approach to tractable choice need to be studied together.

A letter from George Christodoulou, Elias Koutsoupias, and Annamaria Kovacs summarizes their breakthrough proof of the Nisan-Ronen conjecture, which they recently published in STOC'23. Their proof strengthens their prior work on the problem from FOCS'21, which was featured in last year's summer issue of the Exchanges. They discuss the new ideas that helped them finally settle the problem.

Etan Green and Barry Plunkett describe their work applying deep reinforcement learning to eBay bargaining. They nicely summarize the challenges with putting deep RL into practice, as well as the ways they were able to overcome these to outperform human bargaining. Their work received the Best Paper Award at EC'22.

This issue also includes an annotated reading list from Eric Balkanski, Vasilis Gkatzelis, and Xizhi Tan on learning-augmented mechanism design. They consider a mechanism designer with side information about agents' private types. The challenge is to design a mechanism that uses this extra information effectively without giving up worst-case robustness. This list may be of particular interest to readers wishing to follow up on the tutorial on the same subject at EC'22.

This issue ends with a solution to a puzzle by Vincent Conitzer in last year's summer issue on the communication complexity of planning a workshop to celebrate Noam Nisan's 60th birthday, and a new puzzle by Vincent Conitzer on simulation and cooperation in normal-form games in celebration of Joe Halpern's 70th birthday.

We would like to take this opportunity to thank outgoing co-editor-in-chief Inbal Talgam-Cohen for her outstanding service to our community since 2020. We also extend thanks to Yannai Gonczarowski for his continuing help in putting together the issues of Exchanges. As always, please do not hesitate to reach out to us if you would like to volunteer a letter, survey, annotated reading list or position paper. We hope you find the research showcased in this issue inspiring!

Statements from SIGecom Officers Candidates

IRENE LO

Stanford University

and

SAM TAGGART

Oberlin College

SIGecom elects members of its community to various officer positions. Below you will find statements from this year's candidates. All positions have two-year terms, with two-year renewal pending approval of all officers. Members of SIGecom will be notified of voting by email from the ACM. To join the SIG online, visit <https://www.acm.org/special-interest-groups/sigs/sigecom>.

Michal Feldman (Chair). I am deeply honored to be considered for the position of SIGecom Chair. With internet-based systems increasingly governing our economic and social interactions, the work of SIGecom has become increasingly more important and relevant.

Over the past two decades, ACM SIGecom has made significant contributions to research at the intersection of computer science and economics. Looking ahead, it is crucial that we continue to foster an inclusive environment, welcoming different subareas within CS, including theory, AI, and ML, as well as economics, game theory, and operations research. With its unique position at the forefront of this field, ACM SIGecom is ideally positioned to help its members discover new scientific avenues, develop new tools, and train the next generation of researchers and practitioners.

As SIGecom Chair, my goal will be to ensure that our community remains vibrant, welcoming, and focused on tackling the most important challenges and opportunities at the interface of CS and economics. I aim to foster deeper interactions with related areas of science, creating a more collaborative and interdisciplinary environment that will enable us to make even greater strides in this exciting field. To achieve this, it is essential that we build a safe and inclusive community.

Federico Echenique (Vice Chair). My research has dealt with algorithmic questions in economics for many years, and I was one of the first economists to attend, and become involved, in the EC conference. I have been on the program committee for most years since 2012, and served as co-chair (with Shuchi Chawla) for the conference in 2021. If I am elected, I hope to continue the work of the current group of officers, who have overseen an expansion of the EC community, with added diversity of fields and people involved. I am aware that the growth of the SIG presents some challenges, and look forward to thinking about solutions with the new chair and the new slate of officers.

Alexander Teytelboym (Vice Chair). I'm Alex and I'm an economist who has been adopted by the EC community. Although my wonderful "opponent" has an absolute advantage over me on every dimension, I'm nevertheless running for Vice

Chair for three reasons (where I might possibly have some comparative advantage). First, I want to do more to improve the mental health of the members of our community and focus more on the well-being of junior researchers during conferences and the job market. I'm currently a mental health champion and a harassment advisor at the economics department at Oxford. Second, I want to scale up the existing mentorship programme. I've been involved in organising the EC Mentoring workshop and I believe that many students outside elite CS departments are still missing out on the best networking opportunities. Moreover, I think we need to build stronger mentorship programmes for junior researchers approaching tenure. Third, I hope to encourage much more involvement of the community with the third sector (non-profits, NGOs etc). My experience of working with NGOs (on refugee resettlement in particular) has been simultaneously challenging and rewarding, but I think our community has at least as much to add to the third sector as it does to industry collaborations. Thank you for your consideration!

Brendan Lucier (Secretary-Treasurer). It is an honor to be nominated for Secretary-Treasurer of SIGecom. This SIG and the EC conference have been my home base for many years, and I would be thrilled to better serve the EC community in this role.

I believe that SIGecom serves a crucial role in creating opportunities for our members to engage with each other and be recognized for their accomplishments. To that end, I co-chaired the inaugural SIGecom Winter Meeting in 2021. I also led the award committee for the Best Presentation by a Student or Postdoctoral Researcher at EC the first year that award was offered. If elected I will strive to continue the SIG's work of engagement and recognition.

I'm also very familiar with the ACM EC conference itself, having served as Local Arrangements Chair and Conference Treasurer, as Tutorial Co-Chair, and as Theory Track Co-Chair. I am excited to put that experience to use as the executive committee manages EC and other activities in the coming years.

SIGecom is a unique group that brings together researchers across many backgrounds. If elected secretary-treasurer, I will work to strengthen the relationships between and within disciplines that make SIGecom thrive.

Matt Weinberg (Secretary-Treasurer). I'm very honored to be nominated for the SIGecom Secretary-Treasurer position! The EC community has been my research home since I started my PhD, and I'm excited about the opportunity to serve. I've previously served as co-editor for Exchanges, and treasurer for EC, and on the EC Program Committee (as a PC member, SPC member, and AC). I also co-organized this year's SIGecom winter meeting. Finally, I also co-organized the first three EC Mentoring Workshop, and am still active in assisting current organizers. As an officer, I'd be especially excited to continue providing logistical support for similar initiatives, and helping to amplify the vision of other officers and members of the EC community.

SIGecom Winter Meeting 2023 Highlights

EMILY RYU

Cornell University

and

CHENGHAN ZHOU

Princeton University

and

MARYAM BAHRANI

a16z crypto

***Emily Ryu** is a rising third year PhD student in Computer Science at Cornell University, advised by Éva Tardos and Jon Kleinberg. Her research interests span algorithmic game theory, combinatorial optimization, and networks, particularly with more realistic models of behavioral and cognitive constraints. Before Cornell, she graduated from Princeton University with a B.A. in Chemistry and minors in applied math and computer science.*

***Chenghan Zhou** is a rising second year MSE student in Computer Science at Princeton University, advised by Matt Weinberg. Previously, she graduated from University of Virginia with a B.A. in Computer Science, and spent a year visiting Institute for Theoretical Computer Science at Shanghai University of Finance and Economics. Her research interests lie in the intersection of Computer Science and Economics, with a focus on computational economics, analysis and design of algorithms, algorithmic game theory and mechanism design.*

***Maryam Bahrani** is a researcher at a16z crypto, where she studies implications of strategic behavior across layers of the blockchain — from economic security at the consensus layer to efficiency guarantees at the application layer. Before joining a16z, she was a PhD student in the CS theory group at Columbia University, advised by Tim Roughgarden. Prior to that, she completed her undergraduate degree at Princeton, where she worked closely with Matt Weinberg.*

The third annual ACM SIGecom Winter Meeting took place on February 22, 2023. Organized by Scott Kominers and Matt Weinberg, this year’s meeting brought together researchers from economics, computer science, and adjacent fields to focus on Web3, blockchains, and cryptocurrencies. The virtual meeting included talks from and discussions with leading experts on getting into the research space, interesting technical questions, and exciting challenges and opportunities that lie ahead. The day also included interactive exercises that gave participants the opportunity to gain hands-on experience and have fun with NFTs.

Here, we share some highlights and insights from the 2023 Winter Meeting.

Web3: What and why?

The first session of the day was an introduction designed to give everyone the necessary background for the rest of the program, addressing the question: *what* is

blockchain/Web3? The workshop then moved into a more technical panel discussion hoping to answer the question: *why* think about blockchain/Web3?

What: Intro to blockchain and web3

We are increasingly hearing that blockchains are a new and exciting topic at the intersection of economics, computation, and algorithmic game theory. ChatGPT even tells us that Web3 represents “the next evolution of the internet.” But, what exactly *is* a blockchain? To help answer this question, Jacob Leshno began with a presentation intriguingly titled *Blockchain, web3, the promise of decentralization, and soup*. In this primer on blockchains and Web3, Jacob set out to help participants understand the fundamentals of these technologies – what they are, what they are supposed to be, and what they could become.

Jacob suggested that underneath all the hype, expectations, and vague promises to make life amazingly decentralized (and simply *better*), a blockchain is fundamentally an idealized “computer in the sky” – a system to provide trusted storage of data and execution of code, open to all users and not controlled or owned by any single user. This can be implemented using an open write-only ledger, which allows users to commit via automated execution of code (known as “smart contracts”).

So if this magical computer in the sky can be implemented reasonably efficiently, why hasn’t it solved all our problems already?

It turns out that the barriers to decentralization are often not due to the computer itself, but rather the larger legal, political, and social environments in which it operates. Consider the tragic tale of the cryptobros who tried to buy a rare *Dune* book and convert it into NFTs, only to learn that \$3 million could get them a very expensive copy but not the actual rights to the book – since intellectual property is governed by the US courts, not code on a blockchain.¹ And copyright law is just one such instance within a sociopolitical structure that concentrates power and resources in centralized institutions.

What, then, is the secret sauce that blockchain provides to tackle these issues? Or is it just the stone in the proverbial stone soup,² contributing only in name but not in substance to an elaborate mixture that we’ve thrown together?³

Either way, Jacob argues that there are still many reasons to be excited about blockchains. Even if blockchains/Web3 are not the end-all be-all solution, they still draw attention to important legal, economic, and financial questions. Re-examining these systems may then lead to broader impacts on products and markets: For example, how will large financial corporations adapt to a world where individual users can perform basic services on the blockchain on their own? And what implications will this have on preserving competition, openness, and fairness? Last but not least, on the technical side, open decentralized protocols have numerous interesting properties that may have the potential to solve other types of problems as well.

The rest of the presentation explored these points in further depth, using Bitcoin as a running example. Fundamentally, a payment system should store user balances and allow legal transfers. The traditional solution of a centralized operator makes

¹<https://www.esquire.com/entertainment/books/a38815538/dune-crypto-nft-sale-mistake-explained/>

²https://en.wikipedia.org/wiki/Stone_Soup

³And on this note, the promised titular soup was indeed delivered.

this easy, but requires institutional trust, allows for monopolies that can harm welfare, and is vulnerable to hold-up problems. As a decentralized payment system, the Bitcoin protocol aims to address these problems. In contrast to a *platform* with a centralized planner, the term *protocol* suggests a *market-based* approach, where miners provide the infrastructure so that users can obtain services, and anyone can be a miner or a user. Some of the challenges faced by such a decentralized system include validating transfers and ensuring consistency in the presence of nodes that may fail or be malicious, while avoiding control by a single monopolist. By combining tools from cryptography (digital signatures for authentication) and distributed systems (BFT consensus algorithms), Bitcoin’s Nakamoto protocol provides a solution for processing transactions.

To close, Jacob gave a sampling of a few directions for future work. Systems-level questions include optimizing protocols for efficiency, security, and scalability. The design of systems and protocols also inherently involves game-theoretic questions, such as handling collusion between users and miners. Finally, there are broader organizational questions of how blockchain/Web3 can be incorporated into markets, economies, and other structures to create new societal systems to be studied.

With the market-based view of protocols, it is clear that tools from market design will be essential moving forward. Ultimately, Jacob proposed thinking of blockchains as a new paradigm combining elements of open-source software and market design, drawing an active community of researchers from a wide range of fields together to ask and answer many exciting questions.

Why: Panel discussion on interesting questions and challenges

In the panel with Barnabé Monnot, Andrés Monroy-Hernández, and Elaine Shi, we took a deeper dive into some of these exciting questions.

The panelists started off by sharing their paths to the blockchain community. **Barnabé’s** academic roots are in the EC community, having focused on algorithmic game theory and systems during his PhD. A key moment was when he attended Devcon (the Ethereum Foundation’s annual conference) in 2018, where he was so energized by many conversations that he decided he wanted to work full-time for the Ethereum Foundation. **Andrés** comes from a background in Human-Computer Interaction and social computing, and is interested in systems that enable new forms of collaboration. He first started thinking about blockchains at Snapchat, exploring the creation and monetization of digital content on the blockchain. **Elaine’s** background is in systems security and cryptography. She first started working on distributed consensus protocols during the early days of Bitcoin, and became intrigued by how the incentive design of the protocol encouraged early adopters.

Next, each panelist shared one key aspect that they find interesting about the Web3 space. **Barnabé** is interested in the democratization of mechanism design, as blockchains emphasize putting the tools of mechanism design in the hands of users rather than only the platforms. On a similar note, **Andrés** feels that blockchain technologies may transform the underlying infrastructure of online societies by shifting the locus of power from a few to many. **Elaine** noted that she was initially drawn in by the technical challenges, but she has remained an active member of the blockchain community because of its truly interdisciplinary nature, and its unique eagerness to deploy state-of-the-art academic research in the real world.

Each panelist was then asked for one area in which they see blockchains having a tangible impact within the next few years. **Barnabé** noted that impact is not even necessarily a few years away, but is already being made now, for instance, through the huge market cap of Bitcoin and NFTs. **Elaine** highlighted more far-reaching economic implications, such as antitrust regulation and removing monopolies. **Andrés** focused on the popularization of DAOs (decentralized autonomous organizations) as an alternative to incompetent, corrupt, and/or otherwise dysfunctional institutions.

Then, each panelist shared a research direction of interest to the EC community. **Elaine** pointed to decentralized mechanism design – how does the feasibility landscape change with the challenges of a decentralized setting (in which many classical solutions completely fail), but with the help of tools from cryptography? And in addition to incentive compatibility and collusion resistance, what about maximizing welfare, revenue, and other objectives? **Barnabé** highlighted questions about the credibility of auctioneers in decentralized environments, as well as the ecosystem of MEV (miner/maximal extractable value, a broad term referring to the positive and negative externalities arising from multiple users having conflicting goals on shared state machines) – how can MEV be managed, mitigated, captured, and/or utilized? **Andrés** brought up new questions in market design, particularly in three-sided marketplaces (*e.g.*, food delivery systems, with restaurants, drivers, and customers). How can we better design systems with multiple stakeholders and both short- and long-term incentives in mind, and what implications might this have on cooperative ownership, competition, and other dynamics?

To close, the panelists each shared one thing they love most about the blockchain community, one thing they like the least, and concrete suggestions for students and researchers hoping to learn more about the space. **Barnabé** loves how open the community is to new people, feedback, and ideas. **Andrés** loves the new optimism and excitement around the opportunity to start from scratch and reimagine the organizations, platforms, institutions that we have today with an angle of social justice. **Elaine** loves the eagerness to deploy SOTA research, which is a key driving force in making advances in areas ranging from zero knowledge proofs to formal verification and mechanism design. The panelists all expressed frustration at confusing, opaque, and overly financialized systems that have left room for bad actors. In light of this, they emphasized connecting with leaders and mentors for energy, inspiration, and guidance in navigating the complex blockchain space. In particular, **Barnabé** recommended reaching out both at in-person events and online; **Andrés** pointed to resources such as The Blockchain Socialist podcast and SIGCHI papers; and **Elaine** highlighted workshops that bring together academic and industry researchers in an intimate setting (such as this Winter Meeting!).

Overall, the panel concluded on a note of cautious excitement – while blockchains and Web3 may not be the *only* solution to *all* our problems, they are certainly an fascinating path toward exploring technical foundations, socioeconomic dynamics, and other interdisciplinary questions running through our society.

Fireside chat with Tim Roughgarden

In the early afternoon, there was a fireside chat with Tim Roughgarden. Tim is head of research at a16z crypto and Professor of Computer Science at Columbia University. Much of the Q&A centered around three apparent pivots in his career (spoiler: He views them not actually as pivots, but as natural transitions guided by his research interests), followed by his visions for the growth of the blockchain space. Below are some edited excerpts from the Q&A with Tim.

“Pivot” 1: Blockchains and Web3. You were working on game theory and mechanism design for many years. What caused the pivot to blockchains?

That’s a question I get a lot. I can see how from the outside this may seem like a pivot, but from my perspective, it feels like a very natural segue. I’m still doing research in theoretical computer science, which is my core training. When I was a Ph.D. student at Cornell in the late 90s, the internet was blowing up, which drove a lot of important research in computer science. It was also clear that computer scientists needed to learn game theory to reason about applications arising on the internet. Now, as someone who works on foundational computer science in areas that are less well-understood and involve game-theoretic reasoning, blockchains are a very natural application. The economic issues are intertwined with the technology in a more intrinsic way than I’ve ever seen before. There are so many opportunities for computer scientists who do mechanism design and economic theory, and it’s a perfect fit for me as a lifelong EC person, and for this community.

In addition, I would say Web3 is not just a branch of mechanism design, but an entirely new discipline in computer science unfolding in front of our eyes. It’s unbelievably interdisciplinary – drawing from classical ideas in consensus and cryptography and mechanism design. In many ways, doing Web3 research now feels similar to when I worked on algorithmic game theory for the internet in the 2000s. It felt like a new area. Everybody working on it seemed very confused. We made a lot of mistakes. We reinvented the wheel a bunch of times. There were no textbooks or lecture notes written for computer scientists. But on the other hand, if you want to do research that will show up in textbooks, working in an area that doesn’t yet have textbooks is a great way to achieve that.

“Pivot” 2: The application layer. Can you tell us more about your work at the application layer? Automated market makers (AMMs) and decentralized finance (DeFi) seem like a bit of a pivot—what happened there?

The work on DeFi and AMMs also doesn’t feel like as much of a pivot to me. As a theoretician, I’m unusually agnostic about the techniques I use; ultimately, I am more of a problem-driven person. My research agenda is shaped by questions like *what systems and applications do I want to understand? What is the type of math that is appropriate?* Given my focus on Web3, it’s natural to study AMMs.

One of the things that blockchains bring about are questions about the rules of ownership and exchange. In the centralized world, this is well-solved by traditional finance. However, as a “computer in the sky,” the blockchain is quite weak (maybe as powerful as a computer 50-60 years ago), so you have to limit yourself to very simple computations (e.g., order books are prohibitively expensive). A lightweight alternative is automated market makers (AMMs), actually originally developed for

prediction markets to address liquidity issues. In analyzing these alternatives, it was interesting to look at new types of math not necessarily as familiar to the EC community (e.g., continuous-time finance, Black-Scholes, etc.).

“Pivot” 3: a16z and transition between academia and industry. What has working at a16z been like? What are you trying to accomplish there?

When I was asked to start a crypto research lab at a16z, I couldn’t pass on such an incredibly unique opportunity. Researchers at a16z crypto devote about two-thirds of their time to fundamental academic work, with the remaining third collaborating with portfolio companies on super early-stage products, through which it is much easier to have immediate direct impact. Talking with real-world practitioners also reveals fundamental challenges, limitations, and possibilities, helping to identify promising research directions.

What are my hopes for the legacy of it all? I would love if, in hindsight, the a16z lab comes to be viewed as an “inflection point” in two senses: first, as a nudge towards mainstream adoption of crypto and blockchains; and second, as a milestone for Web3 to be viewed as a hard and fascinating area of computer science, and a serious academic discipline.

Let’s say I’m a student and I’m sold! But as you said, it’s hard to find problems because the field moves so fast. What problems do you see right now that you’d love to see more students working on?

First, a few practical points of advice: If you’re starting out, say a first- or second-year Ph.D. student, find a mentor who’s more calibrated to the field than you are. They don’t necessarily have to be your Ph.D. advisor or at your home institution: they could come from industry, or be a more senior grad student. It’s also worth monitoring the literature (e.g., setting arXiv alerts for keywords like MEV). Some papers will be super convincing and exciting! Some will feel like something’s missing, and you can ask yourself what you thought was missing and sit down and write that theorem or paper.

More generally, here are some trends that we need to understand better:

- (1) Macroeconomic effects of mechanism design: Traditionally, the EC community has focused on game theory and microeconomics, but now, a blockchain’s protocol can directly access and manipulate its entire financial ecosystem. What are the consequences of our mechanism design choices on tokenomics, inflation, and other broad economic outcomes?
- (2) Incentives at the L1 layer (the base network): What can be accomplished by fundamental design decisions (e.g., Ethereum’s recent switch from PoW to PoS)? Can the intuition behind these design decisions be supported by theory?
- (3) The application layers: How are the incentive properties of the base network carried through to applications built on this foundation? What properties arise from economic interactions between different layers of the blockchain stack?
- (4) A unified theory of AMMs: Are some AMMs “better” than others? What is the “right” objective function to optimize?
- (5) MEV: Can we develop a standardized vocabulary and theory to describe this very broad phenomenon?

Social activity: Mint your own NFT!

During a 15-minute break, Matt and Scott invited every member of the audience to mint a real NFT of their own. Participants who did not already own a crypto wallet could create one on MetaMask (Matt and Scott provided a direct link to the official website, to avoid scams or phishing websites from search engines), a web browser extension and mobile app that manages users' Ethereum private keys. Creating a wallet involves the generation of a *seed phrase*, or a top secret sequence of words that can be used to access the contents of the wallet. Matt and Scott advised the audience to write down several copies of their seed phrase (by hand!) to store securely in different locations, but *never* to screenshot or type directly into a computer.⁴

After everyone had created their crypto wallet, MetaMask displayed the asset on the Ethereum blockchain (initially 0, since nobody had yet minted a token) with options to buy, send, and swap cryptocurrencies. Then, when participants scanned a provided QR code, they were taken to the POAP website (Proof of Attendance Protocol, a type of NFT). By minting their own POAP, everyone was able to immortalize their attendance at the 2023 SIGecom Winter Meeting to live on the blockchain forever!

NFT case study: Bored Ape Yacht Club

If you could have had the option of buying a Bored Ape NFT in the initial sale at 0.08 ETH, would you have done so? If you had one today (market “floor” price currently around 75 ETH), would you sell it? Rather than leaping at this appealingly massive value increment presented to them, workshop participants generally voiced uncertainty. Why would anyone be willing to spend millions of dollars on an NFT? Just what exactly are they buying, other than a picture on the internet? How could an internet token really be “worth” such an amount of money? And what does it even mean to “own” an NFT? These questions formed a launching point for an interactive discussion led by Scott Kominers, using the Bored Ape Yacht Club (BAYC) NFTs as a lens to demystify this new class of digital deed.

What does it mean to “own a BAYC NFT”? A BAYC NFT is a blockchain record associated with a unique ape image claiming that a crypto wallet is its current owner, which is used to certify ownership. From the viewpoint of a traditional art market, NFT owners are paying for the image of the ape. While it may seem insane to pay millions of dollars for a mere picture, Scott argued that the digital image itself does indeed hold some functional value. Perhaps the foremost are intellectual property (IP) rights and use rights derived from ownership – BAYC holders are granted full commercial usage rights to any of the ape images they own. In addition, the images draw attention and visibility to this exclusive property right, making the abstract concept of ownership more tangible and attractive. For instance, the “Mutant Serum” airdrop allowed owners of Bored Ape images to create new mutant-inspired NFTs, representing a additional level of exclusive membership in the club.

⁴To recover a wallet, the seed phrase is generally entered by selecting each word one at a time from a larger set of words.

However, most participants remained unconvinced that people would pay a small fortune simply to digitally “own” a picture of an ape with anonymous creators. Scott then invited the audience to brainstorm other values that might contribute to the demand for NFTs. One suggestion was that similar to investing in cryptocurrencies, people might believe that NFTs are a good investment that could significantly increase in value in the future. Another idea was that NFTs derive value from the social status that they confer upon their owners.

In addition to these potential factors contributing to the success of BAYC NFTs, Scott highlighted the community built around the Bored Ape collection. Essentially, a BAYC NFT is equivalent to a membership card of the community, which grants holders access to a members-only section of the BAYC official website, private Discord channels, and exclusive events (such as the only Bored Ape “treasure hunt” competition in September 2021 and the in-person “Ape Fest” celebration in November 2021) with the opportunity to hang out with each other. BAYC owners even have a voice in what the project’s funds are used for, providing another avenue for them to feel like they are contributing to the direction of the community. Also worth noting is that, compared to the general crypto community at large, the community of Bored Ape (or any other NFT) holders is much smaller, which encourages quick changes and growth of an active ecosystem around the NFT. As the only way to participate in this ecosystem is by owning a Bored Ape, the value of the NFT is inherently tied to the value that people find in belonging to the community.

These unconventional aspects of NFT markets are made possible by the underlying blockchain infrastructure. By the decentralized nature of blockchain, NFTs significantly reduce the cost to verify the ownership of an asset, building “sturdy” community for NFT holders. Compared to traditional markets, NFTs also provide more liquidity without centralized intermediaries, which lowers the barrier to entering the market or transferring ownership. Further, blockchains provide standardized and public infrastructure layers that reduce the cost of interoperability and portability. With the prevalence of blockchain, people can simply point to their crypto wallets to publish the same content across multiple platforms.

While NFTs are a trendy topic in the crypto community, there is also a lot of doubt surrounding the viability of NFTs. Scott pointed out that a market for an asset cannot exist without a clear definition of ownership. That said, NFTs propose a new class of digital assets that serve as proof of ownership, so they may potentially result in new types of transactions and marketplaces, and ultimately intriguing new questions to explore in market and mechanism design.

Conclusion

The area of Web3, DeFi, and blockchain technology is evolving rapidly. Every day, entrepreneurs and practitioners are building on the theoretical insights from cutting-edge academic research to create innovative new technologies. At the same time, the field can often feel mysterious, even to experts – there is not yet much consensus in the community with respect to basic definitions and questions, not to mention approaches and solutions. Ongoing research is still trying to gain a comprehensive understanding of blockchain technology, and as such, this year’s SIGecom Winter Meeting was largely expository and exploratory. The discussions,

both technical and non-technical, were highly clarifying and inspiring. For computer scientists and economists in the EC community starting to think about blockchain and Web3, the road ahead may be challenging – but this meeting highlighted the many exciting discoveries to be made, and the supportive community of like-minded researchers driving this thriving field.

Report on the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO 2022)

ELENA FALCETTONI

Federal Reserve Board

and

DINA MACHUVE

Nelson Mandela African Institute of Science and Technology

and

BRYAN WILDER

Carnegie Mellon University

and

ANGELA ZHOU

USC Marshall School of Business

1. INTRODUCTION

The second annual ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO'22) was held from October 6-9 at George Mason University in Arlington, VA, USA. This was the first in-person version of the conference: the first event was held virtually in 2021. The conference builds on a line of workshops on Mechanism Design for Social Good (MD4SG) held previously at the ACM Conference on Economics and Computation, and is affiliated with the broader MD4SG initiative. In both 2022 and 2021, SIGecom was a sponsor of the conference. EAAMO aims to highlight work where techniques from algorithms, optimization, and mechanism design, along with insights from the social sciences and humanistic studies, can help improve equity and access to opportunity for historically disadvantaged and underserved communities. A key goal of the conference is to bridge research and practice in this area. Accordingly, we aimed to foster an interdisciplinary community, including researchers from a computer science, operations research, economics, policy, and more, as well as practitioners and policymakers working in areas related to inequality. The conference featured contributed papers and posters as well invited keynote talks, a panel discussion, a doctoral consortium, and community-building activities and social events.

2. KEYNOTE SPEAKERS AND PANELISTS

The conference featured three keynote speakers.

Karen Smilowitz, James N. and Margie M. Krebs Professor in Industrial Engineering and Management Sciences, Northwestern University, gave a talk on *Emerging trends and new research directions in volunteer management*

Karen's talk overviewed a long line of work, including her own in logistics and volunteer management, discussing nonprofit operations and volunteer management. Nonprofit operations have pressing needs, but often rely on a volunteer base and

therefore cannot match supply and demand of labor using wages. Novel opportunities for volunteer engagement, and algorithmic developments, arise with the prevalence of online platforms.

Marcella Alsan, Professor of Public Policy, Harvard Kennedy School, gave a talk on *Representation and extrapolation: Evidence from clinical trials*.

Marcella Alsan is an applied microeconomist studying health inequality. Her talk discussed the consequences and causes of low enrollment of Black patients in clinical trials. In extensive survey experiments, they find that physicians are more willing to prescribe drugs tested in representative samples. They also develop a model of extrapolation in which evidence from representative clinical trials is more likely to affect decision-making. The increased costs of representative enrollment and these benefits of representation can explain the persistence of health inequalities.

Sello Mokwena, Professor of Computer Science, University of Limpopo, gave a talk on *Factors influencing low adoption rate of technologies in developing countries*.

Sello’s talk discussed studies on factors affecting technology adoption in developing countries, using surveys, qualitative analysis, and theories of technological diffusion in information systems. Sello’s work surfaces common themes as to why South African Small Medium Enterprises (SMEs), consumers and government institutions (especially in rural areas) in developing countries face difficulties in adopting technology, such as awareness, cost reduction requirements, and complexity issues.

The conference also featured a panel discussion. The panel included Emanuela Galasso, Senior Economist in the Development Research Group (Poverty and Inequality Team) at the World Bank; Rebecca Johnson, Assistant Professor at McCourt School of Public Policy, Georgetown University; and Sello Mokwena, Professor of Computer Science, University of Limpopo. The panel topic was “Opportunities and Barriers in Bridging Research and Practice” and was moderated by Sera Linardi.

3. CONFERENCE PROGRAM

Conference program overview.

We received over 150 submissions for publication from over 20 countries around the world and across fields, spanning authorship from researchers, policymakers, as well as other domain experts and professionals. All contributors were united by their interest in improving equity and developing solutions for problems in a variety of application domains such as education, labor, environment, healthcare, algorithmic fairness, and digital platforms. Due to its interdisciplinary nature, the conference attracted a very diverse and large group of members with backgrounds in computer science, A.I., operations research, economics, public policy, and humanities, while a great number of papers combined methodologies and insights from multiple fields. Each contributed paper was rigorously peer-reviewed by members of a program committee who were chosen from fields related to the topics of the conference. Out of all submissions, 39 were accepted for oral presentation and 55 were accepted for poster presentation.

From the papers accepted for oral presentation, 19 opted to appear in the second volume of the archival track of the conference: the 2022 Proceedings of EAAMO ‘22: Equity and Access in Algorithms, Mechanisms, and Optimization published by

the ACM. The conference also provides a non-archival presentation option, aiming to enable participation by researchers in journal-focused fields. Out of the accepted papers, we gave awards in the following categories: Best Paper, Best Paper with a Student Presenter, and New Horizons.

Conference program themes.

Our program was organized into paper sessions. Many themes emerged, cross-cutting across methods (theory, algorithms, economics, operations research, data science/machine learning, empirical studies, policy analysis) and application areas. Many themes were related to policy design.

One common thread was a focus on *the empirical, algorithmic, economic and operational* modeling of service delivery, constrained allocation, and incentives. Motivating domains included public sector service delivery (e.g., homelessness services or education) as well as settings related to development. A range of papers developed theoretical or empirical analyses of such settings, including dynamics such as strategic behavior, online interactions, robustness, and inequalities in access. In the policy & practice track, the conference also featured lessons from the field on public-sector deployments in these areas. A related area of focus was that of computational social choice, including extensive empirical and practical work on case studies, algorithmic work on gerrymandering and redrawing congressional or school district boundaries with an equity objective.

Another common theme was centered on the social sciences, whether via the algorithmic/theoretical study of classical social science models or concepts, or thorough empirical social sciences studies of inequality in policy-relevant settings, including normative theory articulating moral foundations for appealing mathematical definitions of fairness. A theme of growing interest this year was *human factors/human-computer interaction* in understanding perception of algorithms in practice as well as comparing practitioner understandings of diversity with proposed algorithmic notions; i.e. studying human aspects of perceptions of inequality as well as mechanisms for the persistence of inequality.

Paper awards.

The Best Paper award winner was:

- Bias, Consistency, and Partisanship in U.S. Asylum Cases: A Machine Learning Analysis of Extraneous Factors in Immigration Court Decisions by Vyoma Raman, Catherine Vera and C.J. Manna

The Best Paper with Student Presenter awards were:

- Improving Access to Housing and Supportive Services for Runaway and Homeless Youth: Reducing Vulnerability to Human Trafficking in New York City by Yaren Bilge Kaya, Kayse Maass, Geri Dimas, Renata Konrad, Andrew Trapp and Meredith Dank
- On Meritocracy in Optimal Set Selection by Thomas Kleine Buening, Meirav Segal, Debabrota Basu, Anne-Marie George and Christos Dimitrakakis

The New Horizons Award that recognizes a paper that pushes the frontiers of AI research was awarded to:

—Dimensions of Diversity in Human Perceptions of Algorithmic Fairness by Nina Grgić-Hlača, Gabriel Lima, Adrian Weller and Elissa M. Redmiles

Community building events on the program.

The program provided several community-building social events, chaired by Lily Xu and Roozbeh Yousefzadeh. Events included breakout sessions organized by *regional and affinity groups, research fields, application areas*. Affinity groups included Queer and Black affinity groups, and regional groups such as Africa and the Middle East, Latin America and the Caribbean, Asia/Pacific, Europe. Research fields included CS theory and CS fairness, law and public policy, economics, and operations research. Application areas included Healthcare, Housing, Education, Environment, Civic Participation, Algorithmic Bias.

The conference also featured a doctoral consortium, organized by Hamsa Bastani and Juba Ziani. This event provided PhD students (primarily drawn from computer science, operations research, and economics) with the opportunity to meet one another on the first day of the main conference and participate in a roundtable discussion on career paths in the interdisciplinary areas spanned by the conference. The event also featured a poster session open to all conference attendees for students to present their work. Finally, students were placed into groups to meet with faculty mentors over the course of the conference. We believe that this program provided an important means for graduate students, especially those new to the event, to become part of the community.

Finally, the conference included a junior faculty session organized by Nikhil Garg and Faidra Monachou. Attendees were primarily junior faculty in computer science, operations research, and economics departments. The session served as the launch for a MD4SG junior faculty network. Participants felt that there would be great value to enabling exchange of experiences, both from senior faculty and between junior faculty, about the process of doing impact-focused research. For example, a common set of questions relate to building effective collaborations with non-profit partners. These reflections have helped set priorities for later junior faculty network events, e.g., a virtual panel on nonprofit collaborations. There was also interest in organizing social events in conferences as well as getting advice on topics such as grant writing, student advising etc. There is especially a lot of junior researcher energy across disciplinary fields, and it would be good to find ways to support/continue to collaborate across such fields.

Reviewing.

We had a two-stage review process. In the first stage, paper received three reviews from an initial set of program committee members, assigned based on bids. At the end of the first stage, reviewers engaged in a discussion process with area chairs to attempt to come to consensus on a decision. For papers where an additional perspective was needed, a new reviewer was assigned in the second stage. These reviewers were hand-selected by the area or program chairs to provide the expertise needed to reach an informed decision on each paper. We found this process to be particularly helpful because of the interdisciplinary nature of the conference. For many papers without consensus at the first stage, reaching a well-informed decision became much easier after adding a reviewer with the right subject-matter expertise,

who may have been in a different field than those assigned in the first stage.

Travel grants and other support.

An important component of the conference organization was providing travel grants to encourage as diverse participation as possible. Financial assistance was provided in the form of: registration waivers, travel grants and accommodation grants. Waivers were provided for forty-five in-person and forty-one virtual registrations, as well as sixteen travel grants and twenty-three accommodation grants. A particularly exciting program, facilitated by Francisco Marmolejo-Cossio, sponsored 7 female indigenous students from Mexico, who presented work from summer research with the Mechanism Design for Social Good (MD4SG) community during a dedicated poster session at the doctoral consortium.

4. CONCLUSION

The cross-cutting themes of conference papers resulted in sessions that featured multiple modes of analysis (for example, quantitative and qualitative) on similar topics. We are excited that the topical focus of EAAMO can bring together different perspectives and facilitate discussions among researchers in different fields studying similar systems and phenomena toward goals of improving equity. Given the breadth of disciplinary perspectives and networks, we thought that topically organized breakout sessions and community-building activities such as the junior faculty network were particularly helpful in highlighting shared interests among attendees.

We believe we succeeded in creating an inclusive in-person conference on how computational tools and algorithms, together with economic approaches and mechanism design, can address equity, access, and other urgent societal challenges.

Acknowledgements

We are grateful to the general chair Sanmay Das and the executive committee, Rediet Abebe, Kira Goldner, Maximilian Kasy, Jon Kleinberg, Illenin Kondo, Sera Linardi, Irene Lo, Ana-Andreea Stoica. We are also extremely grateful to all our committee co-chairs, whose help was absolutely instrumental in putting the conference together: Tasfia Mashiat, Gaurab Pokharel, Lily Xu, Roozbeh Yousefzadeh, Mir Masood Ali, Jose M. Alvarez, Alex DiChristophano, Sandro Radovanović, Rhea Tibrewala, Chinasa Okolo, Sara Kingsley, George Obaido, Juba Ziani, Hamsa Bastani, Gourab Patro, Jessie Finocchiaro, Tasfia Mashiat, Aviva Prins, Francisco Marmolejo-Cossio, Nikhil Garg, and Faidra Monachou. Our breakout sessions were enabled by facilitators: Henry Zhu, Benjamin Laufer, Kenya Andrews, Chinasa T. Okolo, Yaren Bilge Kaya, Francisco Marmolejo Cossío, Niclas Boehmer, Harsh Nisar, Kate Donahue, Alex DiChristofano, Thomas Krendl Gilbert, Shubham Singh, Mir Masood, Edosa Leta, Brian Brubach, Joel Miller, Jiyoo Chang. Thank you to all! We are also grateful to our area chairs who were instrumental throughout the review process: Nikhil Garg, Elisabeth Paulson, Milena Almagro, Neha Bairoliya, Michael Best, Sam Taggart, Eric Chan, Fernando Delgado, Moritz Hardt, Allison Koenecke, Jack Hessel, Matthew Olckers, Alex Teytelboym, Jose Correa, Maria Rodriguez, and Rebecca Wexler.

Finally, we are grateful to our sponsors, the ACM SIGecom and SIGAI, for sponsoring the conference for the second year in a row. We are particularly grateful to our funders, the Sloan Foundation, MacArthur Foundation, and the Public Interest Technology University Network (PIT-UN), whose support has been essential in funding under-represented groups to participate in the conference and in lowering registration fees in general to promote inclusivity. Our funders' support has been used for travel grants for students coming from the Global South, registration fee waivers, and speaker honoraria, which has contributed to a diverse and multi-disciplinary program.

Recent Developments in Pandora’s Box Problem: Variants and Applications

HEDYEH BEYHAGHI

Carnegie Mellon University

and

LINDA CAI

Princeton University

In 1979, Weitzman introduced Pandora’s box problem as a framework for sequential search with costly inspections. Recently, there has been a surge of interest in Pandora’s box problem, particularly among researchers working at the intersection of economics and computation. This survey provides an overview of the recent literature on Pandora’s box problem, including its latest extensions and applications in areas such as market design, decision theory, and machine learning.

1. INTRODUCTION: THE CANONICAL PANDORA’S BOX PROBLEM

In many economic situations, search problems involve multiple options with unknown rewards. Gathering more information can reduce uncertainty about an option’s reward but at a cost. The goal is to obtain a high-quality reward while minimizing the cost. For example, a company seeking to hire job candidates may need to conduct expensive onsite interviews to better assess candidate quality. Similarly, a student choosing between multiple university offers might need to visit campuses to gain a clearer understanding of their preferences for each institution.

The foundational model of optimal search, known as *Pandora’s box problem*, was first established by Weitzman [1979]. The problem consists of a searcher who can choose a prize from one of the n boxes. Each box contains a prize with an unknown value drawn from a distribution known to the searcher a priori. The value distributions of the boxes are independent from each other but may be different. The searcher can perform a sequence of actions, either opening a box or selecting a box. Opening box i incurs a cost c_i , revealing the prize value v_i inside, while selecting box i yields a payoff of v_i and terminates the search process. Importantly, the box must first be opened in order to be selected. The searcher devises an adaptive policy, which determines the next action based on previous actions and outcomes. The objective of the searcher is to maximize their expected utility, which is defined as the expected selected prize value minus the total inspection costs. As an illustration, we use the following running example by Weitzman.

EXAMPLE 1 [WEITZMAN 1979]. *Consider two boxes A and B , where A has a reward of 55 or 100 each with probability 0.5 and a cost of 15, and B has a reward of 0 or 240, with probabilities 0.8 and 0.2, respectively, and a cost of 20.*

Consider a potential strategy for the searcher as follows. The searcher opens box B first. If the reward is 240, the searcher selects the reward and terminates the search. If the reward is 0, the searcher continues on to open box A and takes the

maximum reward observed (here the reward of box A). Using this strategy, for instance, when the reward of the first box is 0 and the reward of the second box is 100, the searcher opens both boxes, paying a total cost of $20 + 15 = 35$, and gains value $\max(0, 100) = 100$ resulting in a utility of $100 - 35 = 65$. Similarly, we can calculate the searcher’s utility in other cases and use the probability of each case to find the expected utility, which for this strategy is 78.

1.1 Optimal Solution and Deferred-Value Interpretation

At first sight, the solution space for Pandora’s box problem seems extremely complicated. In fact, since the optimal policy could be fully adaptive (opening different boxes depending on the history of the boxes that have already been opened and their values), it is not even clear that the optimal policy for Pandora’s box problem can be described in polynomial space as a function of the input size. Surprisingly, Weitzman proves that not only is the optimal solution to Pandora’s box problem efficiently describable, it is highly structured. Specifically, the optimal policy, named as *Pandora’s rule* by Weitzman, is greedy and order non-adaptive (meaning that the inspection order of the boxes is determined apriori, although the policy can adaptively terminate the process). Given any box i , a *reservation value* σ_i can be computed based only on the prize value distribution and the cost for the particular box in question and is not dependent on the value or cost of other boxes.¹ The optimal policy orders the boxes by nonincreasing reservation value and selects the largest observed value once this value exceeds the reservation values of all remaining boxes. Going back to Example 1, as Weitzman shows, although box A has a higher expected value, lower cost, higher minimum value, and lower variance and may seem a better option to try first, surprisingly, it has a lower reservation value and box B will be the first box to open in the optimal solution. Intuitively, by opening box B first, the searcher gains more information about future actions – one can verify that if the searcher opens box A first, the next best action is to open box B regardless of the observed value from box A .

Almost forty years after the introduction of the problem and its optimal solution, Kleinberg et al. [2016] provide a new interpretation of Pandora’s rule that opens the path to new directions in understanding search problems with cost. While Weitzman uses a local improvement argument to prove the optimality of Pandora’s rule, Kleinberg et al. reduce Pandora’s box problem to a related search problem where the values of items are revealed for free. Specifically, Kleinberg et al. define the *deferred-value* of a box i as the minimum between the prize value v_i and the reservation value σ_i . They then prove that the expected maximum deferred value upper bounds the utility of any policy for Pandora’s box problem. Finally, using specific structural properties of Pandora’s rule, they show that the expected utility from Pandora’s rule is exactly the expected maximum deferred value.² As an illustration, in Example 1, the reservation value for boxes A and B are 70 and

¹The reservation value σ_i solves the equation $c_i = \mathbb{E}[\max(v_i - \sigma_i, 0)]$, where the expectation is over the value distribution of box i . The reservation value turns out to be a special case of indices proposed by Gittins [1979]. See Gittins [1979], Weber [1992], and Gittins et al. [2011] for more detail on the Gittins index.

²The reduction of Kleinberg et al. is closely related to the idea of “prevailing charges” in Weber [1992]’s proof of the Gittins index theorem. The deferred value reduction of Kleinberg et al. has

140 respectively, therefore, the deferred value distribution of box A is 55 or 70, each with probability 0.5 and B is 0 with probability 0.8 and 140 with probability 0.2. The expected maximum of these deferred value distributions with no cost is $0.2 \times 140 + 0.8(0.5 \times 55 + 0.5 \times 70) = 78$, which is exactly equal to the maximum expected utility in Example 1.

Since the re-introduction of Pandora’s box problem by Kleinberg et al. to the theoretical computer science community and especially the community working at the intersection of economics and computation, many variants, extensions, and applications of the original model have been considered. For the rest of the survey, we summarize these directions and highlight some of the common themes and techniques.

2. VARIANTS AND EXTENSIONS

In this section, we overview variants and extensions of Pandora’s box problem that have been considered in recent literature. Motivated by the characteristics of specific search applications, these variants and extensions either relax or restrict key aspects of the original model. For instance, there may be multiple ways to inspect a box (Section 2.1), the searcher’s value for the boxes can be correlated (Section 2.4), the cost of inspection may not be additive (Section 2.5), and the searcher may be able to choose more than one item (applied across many variants, often combined with other modifications to the model). In terms of restrictions, the searcher may not be able to inspect in any order they want (Section 2.3), they may not be able to select a previously inspected option that they passed on (Section 2.2), and they may not have exact knowledge about their value distribution for a box (Section 2.4).

Before diving into the specific variants, we will first discuss variations on the objective function and solution concepts that will be used throughout the rest of the section. Firstly, the objective function of the search problems can either be formulated as utility maximization (each box has a non-negative value and the goal is to maximize the selected value minus total cost) or loss minimization (each box has a non-negative price and the goal is to minimize the selected price plus total cost), and both objectives have been studied since the inception of the optimal search problem [DeGroot 1970, Chapter 3]. For exact optimization, the two objectives are equivalent; however, approximating the optimal loss is often easier than approximating the optimal utility. When discussing the variants, we consider the more commonly used utility maximization objective as the default objective, except when otherwise explicitly stated.

In terms of solution concepts, it may not be possible to efficiently describe or compute the optimal policy among all possible policies for certain variants. To overcome this, it is helpful to focus on more limited classes of policies with better descriptive or computational properties. The three most commonly considered solution concepts are as follows:

— **Fully adaptive policy:** the most general class of policies, where the action of

antecedents in Armstrong and Vickers [2015] and has been independently discovered by Armstrong [2017] and Choi et al. [2018].

the policy can depend on previous actions and the values it has seen.

- **Order non-adaptive policy:** the class of policies where the inspection order of the boxes is predetermined before the value of any box is revealed. However, the stopping rule, i.e., when the policy terminates the process, may be adaptive.
- **Fully non-adaptive policy:** the class of policies where both the inspection order and the stopping rule are non-adaptive. In particular, the policy would always inspect all boxes that are specified in the inspection order.

2.1 Alternative Inspection Methods

In many search applications such as student choosing universities or consumer search, there may be several different inspection methods (e.g., online research vs. visiting in person), and inspection may not be required before selecting a box. Models in this subsection relax the original model to allow such variations in methods of inspection.

The most well-studied thread under this relaxation is the *nonobligatory* inspection model. In this model, instead of having to inspect before selecting a box, the searcher can alternatively claim the box closed without inspection and get the expected value of the box. This model has been introduced independently in different communities (wireless network, stochastic testing, search theory) and under different names [Guha et al. 2008; Chang and Liu 2009; Attias et al. 2017; Doval 2018]. In particular, Doval [2018] formulated the nonobligatory inspection model explicitly as a generalization to the original Weitzman’s Pandora’s box problem and popularized the model in the economics and computation community. Recently, a steady line of work [Guha et al. 2008; Doval 2018; Beyhaghi and Kleinberg 2019; Fu et al. 2023; Beyhaghi and Cai 2023] resolved both the computational complexity and approximability of the problem.

The literatures on complexity, structure and approximability of the non-obligatory inspection model progressed in conjunction and are deeply intertwined. Guha et al. [2008]³ first show a significant structural result: the optimal policy claims a unique box closed across all decision branches. Using their structural result, Guha et al. show that the competitive ratio of *committing policies* (order non-adaptive policies where the searcher commits ahead of time to whether they will inspect each box prior to selecting it) is exactly 0.8. Independently, Beyhaghi and Kleinberg [2019] show that the competitive ratio of committing policies is at least $1 - 1/e \approx 0.63$ by a reduction to stochastic submodular maximization, which can also be applied to more general models as we will discuss later.

Doval [2018] provides evidence both for the complexity of the optimal policy and the existence of additional structure. In particular, Doval shows that the optimal policy may be order-adaptive, while showing that the optimal policy has a two-phased structure (where the inspection order only changes once) under additional assumption on the value distribution.⁴

³Guha et al. [2008] studied the problem under the context of stochastic probing in wireless networks. The wider community was unaware of their work until recently.

⁴Specifically, Doval [2018] considers the binary prize environment, where the value of each box i is supported on $\{L, H_i\}$, where the low value in the support is shared between all boxes, but the high value in the support may be distinct for each box.

Finally, Fu et al. [2023] prove that the problem is NP-hard, which confirms intuition in previous literature. In terms of structural results, Fu et al. [2023] and Beyhaghi and Cai [2023] show that in the *general* setting, the optimal policy is two-phased and can be fully specified through an initial inspection order and a threshold for each box.⁵ As a consequence, the decision version of the problem is in NP (since one can prove the optimal utility is above a certain threshold by succinctly describing the policy that obtains this utility). Further, Fu et al. and Beyhaghi and Cai provide a PTAS for the nonobligatory inspection model.

In the nonobligatory inspection model, there are two ways of inspecting a box: pay the full cost and inspect the box or claim the box closed without inspection. However, there may be other options that lie in between: perhaps a smaller cost is needed to reduce the variance of the value distribution. The remaining variants in this subsection addresses the “in between” inspections. Kleinberg et al. [2016] consider an alternative model where there are multiple stages of inspection, and the searcher could only claim the value in the box after all stages of inspections are completed. As the searcher progresses through the stages, more information about the value is revealed, and more cost is incurred. The searcher can stop examining the model further at any stage. Kleinberg et al. find that the optimal policy for this staged inspection setting is a generalized form of Pandora's rule, where a reservation value can be computed for a box at any stage, and the searcher always inspects the box with the highest reservation value (given its current stage).

In a similar spirit, Ke and Villas-Boas [2019] consider a model where information is revealed gradually, but the searcher can claim the box (or stop) at any point. In addition, in their model, the discovery process is continuous rather than discrete, and the value of each box is binary supported.⁶ They find that even in the case when there are two boxes and a fixed-valued outside option, the solution space may be complicated; they characterize the optimal policy under conditions such as when the outside option is below or above certain thresholds.

Aouad et al. [2020] introduce a model where the searcher has two ways of opening a box: fully open and partially open. Similar to Kleinberg et al. [2016], a box must be fully opened before the value can be claimed. However, unlike the model in Kleinberg et al., the searcher can fully open a box without partially opening the box first. Aouad et al. prove that the best *committing policy*⁷ is $(1 - 1/e)$ -competitive to the optimal utility using an analysis inspired by Beyhaghi and Kleinberg [2019]. Moreover, they show that any committing policy or its negation (flipping which box should be partially opened versus not partially opened) is $1/2$ -competitive to the optimal utility. Aouad et al. also design a simple threshold-based committing policy that is near optimal when the number of items is sufficiently large.

Finally, as a direct extension to the non-obligatory inspection model, Beyhaghi [2019] introduces *Pandora's box problem with alternate inspection* model, where the searcher has k different methods for inspecting each box (including not inspecting

⁵The results in Fu et al. [2023] and Beyhaghi and Cai [2023] build on the structural results in Guha et al. [2008] and [Doval 2018].

⁶Note that the support for different boxes could be different, but the value of each box only has two possibilities.

⁷As defined in our discussion of the non-obligatory inspection model.

at all), and the searcher can select at most one method for each box. [Beyhaghi and Kleinberg 2019; Beyhaghi 2019] prove that *committing policies*⁷ are $(1 - 1/e)$ -competitive in this model.

2.2 Search Without Recall

Even before Weitzman’s seminal paper, the economics literature studied both the optimal search problem with recall (the searcher can select alternatives that they have seen in the past) and without recall (the searcher has to select the item or forgo it forever) [DeGroot 1970, Chapter 3]. Hybrid settings where recall is uncertain have also been considered [Karni and Schwartz 1977]. In modern literature, the version of Weitzman’s Pandora’s box problem without recall is studied under the name *Committed Pandora’s box* [Fu et al. 2018].

Kleinberg et al. [2016] first showed that a simple threshold-based policy for Committed Pandora’s box with arbitrary (possibly adversarial) constraints on inspection order achieves at least $1/2$ of the optimal utility by a reduction to prophet inequalities (an online selection model that is similar to Committed Pandora’s box, but without inspection costs). The connection between (variants of) Committed Pandora’s box and (variants of) prophet inequalities is repeatedly exploited by subsequent work.⁸

Fu et al. [2018] and Segev and Singla [2021] provide a general framework for deriving polynomial time approximation schemes (PTAS) and efficient polynomial time approximation schemes (EPTAS) for stochastic optimization problems, respectively. As an application of their framework, Fu et al. show that Committed Pandora’s box problem with free order selection (the searcher has the full freedom to pick the inspection order) has a PTAS by a direct reduction. Segev and Singla show that, in fact, Committed Pandora’s box problem has an EPTAS by first reducing the problem to free order prophet inequality,⁸ which has an order non-adaptive optimal policy, and then applying their framework.

Esfandiari et al. [2019] study Committed Pandora’s box problem under adversarial order and where the searcher is allowed to collect multiple prizes subject to general feasibility constraints (e.g., cardinality, knapsack or matroids constraints). In addition, in their model, the prize values and costs are drawn from a joint distribution, and the cost is only revealed after opening the box. Esfandiari et al. prove that all variants of Committed Pandora’s box problem they consider can be reduced to a corresponding prophet inequality problem with known competitive ratios. Further, they extend Committed Pandora’s box with adversarial order (with the objective of selecting one prize per round) to the contextual bandit setting and obtain a $1/2$ -competitive policy based on the reduction to prophet inequalities.

2.3 Restricted Order of Inspection

As we have discussed in Section 2.2, both free order (no restriction) and adversarial order (complete restriction) are standard assumptions for online selection problems and have been considered in the context of Pandora’s box problem. Motivated by applications such as funding research development, a more general question can be

⁸ See Lucier [2017] and Correa et al. [2019] for in-depth discussions on prophet inequalities and variants.

asked: what if there are *some* restrictions on the searcher's inspection order?

Boodaghians et al. [2020] initiate the study of partial order constraints such as tree or DAG like order restrictions for Pandora's box problem. In the case of (single selection) Pandora's box problem with tree or forest like order constraints, where a box can only be opened once its parent box is opened, Boodaghians et al. show that an order-dependent version of the reservation value can be computed for each box, and the optimal policy always opens the remaining box (if any) with the largest reservation value.⁹ On the other hand, for variants where multiple selections are allowed (e.g., under matroid feasibility constraints or more general constraints), or when the order constraint is DAG like (where a box can be opened only if one of its predecessor boxes is opened), Boodaghians et al. show that finding a fully adaptive policy that achieves ϵ fraction of the optimal utility is NP-hard. They also propose a relaxed definition of approximately optimal policies and analyze the adaptivity gap between fully adaptive policies and fully non-adaptive policies.

2.4 Beyond Independence and Distribution Assumptions

This section considers relaxations on two of the constraints in the original model: the independence of the searcher's valuation among boxes, and the knowledge of the distributions (or even having sample access to the distributions). Since Chawla et al. [2020] introduced both relaxations in the same paper, the study of these variations is often interweaved and thus presented here in a single section. Depending on whether the valuations are independent or correlated and whether historical samples are available (distributional learning setting) or the search is a repeated process without historical samples (online learning setting), there are four different variants that are discussed in this section.

Chawla et al. [2020] study the correlated value and distributional learning setting, mainly under the loss minimizing objective. Specifically, the prices in the boxes are drawn from an arbitrarily correlated joint distribution; moreover, the searcher is limited to $\text{poly}(n)$ samples from the joint distribution, where n is the number of boxes. Chawla et al. show that approximating the loss of the optimal fully adaptive policy within any sublinear factor requires exponential samples. Then, they efficiently find an order non-adaptive policy that is a constant approximation to the best order non-adaptive policy. Moreover, unless $P = NP$, one cannot efficiently find a fully adaptive policy that exceeds the best order non-adaptive policy. Chawla et al. also show that if their model has the utility maximization objective instead, no computationally efficient fully adaptive policy can even be a constant approximation to the best order non-adaptive policy. As a direct follow-up to Chawla et al., Gergatsouli and Tzamos [2023] show that a generalized version of Weitzman's policy is constant-competitive against the best order non-adaptive policy. Compared to the linear programming rounding approach in Chawla et al., Gergatsouli and Tzamos's construction of a competitive order non-adaptive policy is more explicit while obtaining an improved competitive ratio.

Gergatsouli and Tzamos [2022] study the correlated value model under the more

⁹Interestingly, the authors mention that Pandora's box problem with forest like order constraint is closely related to the branching bandit process studied in [Weiss 1988; Keller and Oldale 2003], and their optimal policy has a similar structure.

restrictive setting of online learning and show that when the prices of the boxes are selected by an oblivious adversary in each round, results of Chawla et al. extend to both the full information setting (values of all the boxes are revealed after each round) and bandit setting (only values of the boxes that the policy opened are observable).

Guo et al. [2021] study the independent valuation and distributional learning setting, and prove that the searcher could obtain an ϵ -additive approximation to the optimal utility with high probability given $\tilde{O}(\frac{n^3}{\epsilon^3})$ samples. Fu and Lin [2020]¹⁰ improve this sample complexity to $\tilde{O}(\frac{n}{\epsilon^2})$. Finally, Atsidakou et al. [2022] and Gatmiry et al. [2022] study the independent valuation and online learning setting. Atsidakou et al. extend the original Pandora’s box model (with loss minimizing objective) to the contextual bandit setting, where each round comes with potentially different sets of boxes. At the beginning of each round, the context and cost of the boxes are revealed up front, while the distribution of values in the boxes remains unknown (and can be different from round to round). Atsidakou et al. prove that as long as the context can be used to estimate the reservation value of the box, sub-linear regret against the optimal policy with full distributional knowledge (i.e., Pandora’s rule) is achievable for both the full information and bandit setting. Gatmiry et al. [2022] study the online Pandora’s box problem (with utility maximizing objective) where the boxes have unchanging prize value distributions that are unknown to the searcher. The value distributions are independent but not necessarily identical. Gatmiry et al. prove that in the bandit setting, the searcher can achieve $O(\text{poly}(n)\sqrt{T})$ regret by an algorithm that estimates and then shrinks a confidence interval on each box’s reservation value, where T is the number of time steps.

2.5 Beyond Additive Costs

For many applications that motivate Pandora’s box problem, such as students selecting universities and job search, the cost of inspection may not be additive. For instance, students who visit universities in nearby locations back to back may experience lower costs compared to having three separate trips to those universities. Berger et al. [2023] extend Weitzman’s Pandora’s box problem by considering more general classes of cost functions such as submodular, XOS, or sub-additive functions. Their main result shows that the optimal policy for Pandora’s box problem is order non-adaptive for submodular cost functions. On the other hand, when the cost function is XOS or sub-additive, adaptivity is required for the optimal policy. They also show that even for the more restrictive class of submodular cost functions, approximating the utility of Pandora’s box problem requires an exponential number of queries to the cost function.

3. APPLICATIONS

In this section, we overview applications of Pandora’s box problem in combinatorial optimization (Section 3.1), mechanism design (Section 3.2), delegation (Section 3.3) and matching markets (Section 3.4). The elegant structure of optimal or

¹⁰Although published out of order, Guo et al. [2021] preceeds Fu and Lin [2020] and is cited as prior work in all versions of Fu and Lin [2020].

approximately-optimal solutions to Pandora's box problem plays a crucial role in addressing domain-specific problems where information acquisition is costly.

3.1 Combinatorial Optimization

Kleinberg et al. [2016] and Singla [2018] applied the structures and tools from Pandora's box problem in a wider range of combinatorial optimization problems, such as the costly counterpart of maximum weighted matching, maximum knapsack, minimum vertex cover, minimum set cover, minimum facility location, and minimum prize-collecting Steiner tree. Kleinberg et al. initiated this thread by applying their no-cost reduction explained in Section 1.1. Later, Singla [2018] expanded upon this reduction technique and applied it to a broader range of problems. Singla provides a general transformation for converting *frugal algorithms* (a type of greedy algorithm) into policies for solving combinatorial counterparts of Pandora's box problem, where the searcher can choose multiple boxes subject to feasibility constraints on the selected set. This transformation applies in both utility maximization and loss minimization settings.

3.2 Mechanism Design

We overview several mechanism design papers with costly information acquisition that utilize the structure of optimal or approximately optimal solutions to Pandora's box problem. These papers consider a few different scenarios between sellers and buyers, with a costly investigation process on one side of the market or the other.

Cr  mer et al. [2007] consider an auction scenario for selling a single item, where the set of buyers is not exogenous or determined in advance, and the seller needs to go through a costly sequential process to inform potential buyers about the auction. They show that in the case of independent buyers' valuations, the seller's problem can be reduced to Weitzman's model, where the distribution of each box is the Myerson virtual value distribution for each buyer.

Kleinberg et al. [2016], also focus primarily on the sale of an item to buyers. However, unlike the previous scenario, the buyers are informed about the auction but need to incur inspection costs to determine their values. Using their reduction from costly information acquisition to optimization with no cost (discussed in Section 1.1) as a key element, they devise a descending price auction that achieves the same efficiency as a first price auction with modified value distributions but no cost of inspection, resulting in a small price of anarchy (approximate optimality). Later, Alaei et al. [2021] extend the revenue maximization setting of Kleinberg et al. to the nonobligatory inspection model. They provide mechanisms, both for selling a single item and multiple copies of an item, that are approximately optimal even when the buyers arrive in an adversarial order. Subsequently, Wu et al. [2022] also consider revenue maximization in a nonobligatory inspection setting; however, they particularly focus on the role of bundling the items in optimizing revenue. They study two different markets; the first with one mature and one new product, and the other with two new products. The valuation uncertainty only exists for new products. They show that in a market with one mature and one new product, bundling encourages search, while in a market with two new products, it discourages search. Orthogonally, Fu and Lin [2020] use the correspondence between the descending price auction in a costly information acquisition setting and the first price auction

in the classic setting, developed by Kleinberg et al., to provide sample complexity bounds for auction design with costly inspections.

Armstrong [2017] examines a scenario in which buyers intend to purchase a product, such as a book from an online marketplace, from one of several available sellers. A buyer learns their value and the price of a product upon inspection unless the seller advertises their price, in which case the price is known a priori. When a buyer purchases a product, they get utility equal to their value for the product minus the price and inspection cost set by the seller. From a buyer's perspective, their search problem is exactly equivalent to the canonical Pandora's box problem, and they are modeled to be employing Pandora's rule. Consequently, the reduction of Kleinberg et al. [2016] and Armstrong and Vickers [2015] also applies to Armstrong's setting with prices and can be used to calculate a buyer's expected utility. Armstrong instead focuses on analyzing the seller's strategy decisions (setting product prices, using advertising to guide consumer searches, and determining the consumer's search costs) and their implications, both when a monopolist seller owns multiple products and when there are multiple sellers. Armstrong presents a detailed discussion of the factors that influence which sellers raise or lower their prices given the buyer's inspection order. In addition, Armstrong explores why it might be profitable for a seller to *obfuscate* the searcher by increasing its own inspection cost and examines the equilibria of the buyer-seller optimization problem.

Choi et al. [2018] examine a pricing game where a group of sellers with substitutable items compete with each other while buyers have partial information about their values. In the first step, the sellers simultaneously announce their prices. In the second step, the buyers go through a costly search process among sellers depending on the announced prices and their partial information about their values. The authors characterize buyers' optimal behavior and analyze the pricing game among the sellers.

Chen et al. [2022] propose a three-step mechanism for manufacturers outsourcing their production to suppliers to reduce procurement costs. First, suppliers submit price bids for contracts. Second, buyers investigate ways to reduce production costs, subject to a limit on the number of investigations. Third, the buyer awards the contract to the supplier with the lowest updated bid. The second step, which is a costly investigation process, is equivalent to a variant of Pandora's box problem where there is a limit on the number of boxes that can be opened. Although Weitzman shows that, generally, Pandora's rule may not be optimal given the limitation on the number of inspections, Chen et al. identify sufficient conditions for Pandora's rule to be optimal for buyer investigation, in which case the structural properties of Pandora's rule can be used to design the optimal three-step mechanism.

3.3 Delegated Search

Delegation in search problems refers to the process of a *principal* assigning a search problem to an *agent*, who possesses the necessary resources but may have interests that differ from those of the principal. A key question when considering a delegated search problem is how much the principal loses when they delegate the search to an agent. This quantity is referred to as the *delegation gap*. Although the theory of delegation was introduced in economics much earlier by the work of Holmstrom [1978] and Holmstrom [1984], one of the early papers that use ideas from Pandora's

box problem to design optimal delegated search mechanisms is by Postl [2004] who establishes a condition that ensures there is no loss in the delegation for the same-cost two-box version of the problem. Later, Kleinberg and Kleinberg [2018] use a model introduced in Armstrong and Vickers [2010] and incorporate ideas from Pandora's box problem to create nearly optimal delegated search mechanisms. One of the models they study is a costly information acquisition delegated search problem with binary options. This model involves a set of options, each with a probability of being feasible for the principal and a cost to investigate. They design a search mechanism with a limited delegation gap. Bechtel et al. [2022] later expand on the binary case to include matroid feasibility constraints. However, they prove that there is no constant-factor delegation gap beyond the binary model. To overcome this, they explore other variations, such as the shared-cost model, where the principal can choose how to split the costs with the agent before the delegation. They demonstrate that the shared-cost model has a constant-factor delegation gap for specific constraints.

3.4 Matching Markets

Immorlica et al. [2020] explore a generalization of Pandora's box model within the context of matching markets, specifically focusing on many-to-one markets such as student-college mappings. In these scenarios, students must undergo a costly information acquisition process to determine their values for each college, while colleges maintain a publicly known ranking system for students. The authors introduce *regret-free stability* as a refined solution concept that builds upon the traditional *stability* definition in matching market literature, ensuring optimal information acquisition for students, and they demonstrate the existence of such a solution.

In a single-student model, the problem simplifies to the original Pandora's box problem, making Pandora's rule the optimal solution for the student's search. However, when multiple students are involved, the available college options for each student depend on the valuations of their peers. This interdependency between students' information acquisition choices is resolved by using approximate cutoffs (i.e., the lowest admissible student ranking for each college). With these cutoffs, students can independently tackle Pandora's box problem for the set of colleges where their rank meets the cutoff, ultimately achieving regret-free stability.

4. DISCUSSION: COMMON STRUCTURAL AND TECHNICAL THEMES

Although the variants and applications discussed in Sections 2 and 3 often extend or utilize Pandora's box problem in orthogonal directions, several concepts and ideas appear to be relevant across numerous variants and applications.

In the study of Pandora's box problems, a common theme is analyzing the relative power of simple policies, which typically refers to order non-adaptive policies, and comparing them to fully adaptive ones. This comparison is similar to the concept of the *adaptivity gap* in combinatorial optimization. The variants discussed in Section 2 can be categorized into three classes: order non-adaptive policies being as powerful as fully adaptive policies, having a constant competitive ratio, and exhibiting a super-constant gap between them. The original Pandora's box problem and some models with restricted order of inspection (Section 2.3), beyond additive cost (Section 2.5), and search without recall (Section 2.2) possess optimal order

non-adaptive policies, placing them in the first class. In contrast, various models involving alternative inspection methods (Section 2.1) may not have optimal order non-adaptive policies, but they are constant-competitive against the best fully adaptive policy, falling into the second category. Additionally, the optimal policy in these cases may require limited adaptivity. More sophisticated models with restricted order of inspection (Section 2.3) and correlated distribution (Section 2.4) belong to the third class, as they exhibit a super-constant utility gap between order non-adaptive and fully adaptive policies.

In essence, for different Pandora’s box problem variants, the adaptivity gap serves as an indicator of the problem’s structural complexity. A large adaptivity gap, combined with an impossibility result in approximating the optimal fully adaptive policy, can motivate researchers to focus on approximating the optimal order non-adaptive or fully non-adaptive policy instead.

Another common technical theme is the reduction of a variant of Pandora’s box problem with cost to a related problem without cost. This reduction is most direct in Pandora’s box problem without recall setting (Section 2.2), where many variants can be reduced to different variants of the prophet inequality problem. In the latter problem, the searcher’s values for the boxes are drawn from known distributions and must select a box without recall, but revealing the value does not come at a cost. For the original Pandora’s box problem with recall, Kleinberg et al. [2016] first used a cost-to-no-cost reduction in their alternative proof of optimality for Pandora’s rule. Interestingly, a similar reduction can also be applied to Pandora’s box problem with alternative inspection methods (Section 2.1) and the more general combinatorial optimization setting (Section 3.1). The presence of cost in Pandora’s box problem complicates the design of optimal or approximately optimal policies (approximating values and costs separately may not lead to approximately-optimal utility). Consequently, the reduction from cost to no cost often serves as a useful tool in revealing the structure of the optimal policy or facilitating the design of an approximately optimal policy.

5. CONCLUSION AND FUTURE DIRECTIONS

The recent growing body of research on Pandora’s box problem, presented in this survey paper, has introduced numerous variations and applications, indicating substantial potential for future exploration and investigation.

The alternative and generalized models of Pandora’s box overviewed in Section 2 are far from exhaustive (especially those related to order restriction (Section 2.3), correlated value distributions (Section 2.4), and non-additive cost functions (Section 2.5)), and warrant further investigation. In addition, the relationship between Pandora’s box variants and broader stochastic optimization problems (e.g., Markov chains, multi-armed bandits) has been noted in multiple studies and can benefit from a systematic analysis. We also anticipate the emergence of future models that deviate from existing variants, exploring different aspects of the searcher’s decision-making. For example, current models do not incorporate behavioral economic findings, such as risk and loss aversion. Empirical evidence from Bhatia et al. [2021] demonstrates that these factors align more closely with human decision-making behavior in costly information acquisition settings. In another potential variant,

the searcher may have a long time horizon with the boxes emerging and disappearing in an online fashion. Alternatively, there may be random signals from global events (e.g., the emergence of new technology) that provide information for all or a large segment of the boxes (e.g., a readjustment to the skills of applicants or their distribution).

In the application domain, Pandora’s box model is relevant to most applications where the cost of information acquisition is significant, including those that are not mentioned in Section 3, such as voting and advertisement. In voting scenarios, both the candidates and the voters may engage in costly investigations; e.g., the candidates optimize their investment of targeted campaigning to select populations, while the voters may face a similar trade-off as in canonical examples of Pandora’s box model, where they choose between selecting well-known candidates and investigating their preference alignment with the less well-known ones. Similarly, in advertisement applications, the advertisers go through a costly investigation to select what populations to target and what advertisement methods to use. Other potential applications include models of labor, product, and financial markets with costly information acquisition. In addition, most existing applications employ the canonical Pandora’s box model. Recent work on the variants proposed in Section 2 offers a wider range of modeling choices, and may enable greater realism for some applications.

6. ACKNOWLEDGEMENT

The authors are grateful to Laura Doval, Brendan Lucier, and the editors-in-chiefs Irene Lo and Sam Taggart for their valuable feedback on earlier versions of the survey and identification of missing relevant work.

REFERENCES

- ALAEI, S., MAKHDOUNI, A., AND MALEKIAN, A. 2021. Revenue maximization under unknown private values with non-obligatory inspection. In *EC ’21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, P. Biró, S. Chawla, and F. Echenique, Eds. ACM, 27–28.
- AOUAD, A., JI, J., AND SHAPOSHNIK, Y. 2020. The pandora’s box problem with sequential inspections. *Available at SSRN 3726167*.
- ARMSTRONG, M. 2017. Ordered consumer search. *Journal of the European Economic Association* 15, 5, 989–1024.
- ARMSTRONG, M. AND VICKERS, J. 2010. A model of delegated project choice. *Econometrica* 78, 1, 213–244.
- ARMSTRONG, M. AND VICKERS, J. 2015. Which demand systems can be generated by discrete choice? *Journal of Economic Theory* 158, 293–307.
- ATSIDAKOU, A., CARAMANIS, C., GERGATSOULI, E., PAPADIGENOPOULOS, O., AND TZAMOS, C. 2022. Contextual pandora’s box. *CoRR abs/2205.13114*.
- ATTIAS, C., KRAUTHGAMER, R., LEVI, R., AND SHAPOSHNIK, Y. 2017. Stochastic selection problems with testing. *Available at SSRN 3076956*.
- BECHTEL, C., DUGHMI, S., AND PATEL, N. 2022. Delegated pandora’s box. In *EC ’22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, D. M. Pennock, I. Segal, and S. Seuken, Eds. ACM, 666–693.
- BERGER, B., EZRA, T., FELDMAN, M., AND FUSCO, F. 2023. Pandora’s problem with combinatorial cost. *CoRR abs/2303.01078*.
- BEYHAGHI, H. 2019. *Approximately-optimal Mechanisms in Auction Design, Search Theory, and Matching Markets*. Cornell University.

- BEYHAGHI, H. AND CAI, L. 2023. Pandora’s problem with nonobligatory inspection: Optimal structure and a PTAS. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, B. Saha and R. A. Servedio, Eds. ACM, 803–816.
- BEYHAGHI, H. AND KLEINBERG, R. 2019. Pandora’s problem with nonobligatory inspection. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, A. Karlin, N. Immorlica, and R. Johari, Eds. ACM, 131–132.
- BHATIA, S., HE, L., ZHAO, W. J., AND ANALYTIS, P. P. 2021. Cognitive models of optimal sequential search with recall. *Cognition* 210, 104595.
- BOODAGHIANS, S., FUSCO, F., LAZOS, P., AND LEONARDI, S. 2020. Pandora’s box problem with order constraints. In *EC ’20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, P. Biró, J. D. Hartline, M. Ostrovsky, and A. D. Procaccia, Eds. ACM, 439–458.
- CHANG, N. B. AND LIU, M. 2009. Optimal channel probing and transmission scheduling for opportunistic spectrum access. *IEEE/ACM Trans. Netw.* 17, 6, 1805–1818.
- CHAWLA, S., GERGATSOULI, E., TENG, Y., TZAMOS, C., AND ZHANG, R. 2020. Pandora’s box with correlations: Learning and approximation. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, S. Irani, Ed. IEEE, 1214–1225.
- CHEN, Q., BEIL, D. R., AND DUENYAS, I. 2022. Procurement mechanisms with post-auction pre-award cost-reduction investigations. *Operations Research* 70, 6, 3054–3075.
- CHOI, M., DAI, A. Y., AND KIM, K. 2018. Consumer search and price competition. *Econometrica* 86, 4, 1257–1281.
- CORREA, J., FONCEA, P., HOEKSMAN, R., OOSTERWIJK, T., AND VREDEVELD, T. 2019. Recent developments in prophet inequalities. *ACM SIGecom Exchanges* 17, 1, 61–70.
- CRÉMER, J., SPIEGEL, Y., AND ZHENG, C. Z. 2007. Optimal search auctions. *Journal of Economic Theory* 134, 1, 226–248.
- DEGROOT, M. H. 2004; 1970. *Optimal statistical decisions*. John Wiley & Sons.
- DOVAL, L. 2018. Whether or not to open pandora’s box. *Journal of Economic Theory* 175, 127–158.
- ESFANDIARI, H., HAJIAGHAYI, M., LUCIER, B., AND MITZENMACHER, M. 2019. Online pandora’s boxes and bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. 1885–1892.
- FU, H., LI, J., AND LIU, D. 2023. Pandora box problem with nonobligatory inspection: Hardness and approximation scheme. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, B. Saha and R. A. Servedio, Eds. ACM, 789–802.
- FU, H., LI, J., AND XU, P. 2018. A PTAS for a class of stochastic dynamic programs. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, I. Chatzigiannakis, C. Kaklamanis, D. Marx, and D. Sannella, Eds. LIPIcs, vol. 107. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 56:1–56:14.
- FU, H. AND LIN, T. 2020. Learning utilities and equilibria in non-truthful auctions. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds.
- GATMIRY, K., KESSELHEIM, T., SINGLA, S., AND WANG, Y. 2022. Bandit algorithms for prophet inequality and pandora’s box. *CoRR abs/2211.08586*.
- GERGATSOULI, E. AND TZAMOS, C. 2022. Online learning for min sum set cover and pandora’s box. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvári, G. Niu, and S. Sabato, Eds. Proceedings of Machine Learning Research, vol. 162. PMLR, 7382–7403.
- GERGATSOULI, E. AND TZAMOS, C. 2023. Weitzman’s rule for pandora’s box with correlations. *arXiv preprint arXiv:2301.13534*.

- GITTINS, J., GLAZEBROOK, K., AND WEBER, R. 2011. *Multi-armed bandit allocation indices*. John Wiley & Sons.
- GITTINS, J. C. 1979. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B (Methodological)* 41, 2, 148–164.
- GUHA, S., MUNAGALA, K., AND SARKAR, S. 2008. Information acquisition and exploitation in multichannel wireless networks. *CoRR abs/0804.1724*.
- GUO, C., HUANG, Z., TANG, Z. G., AND ZHANG, X. 2021. Generalizing complex hypotheses on product distributions: Auctions, prophet inequalities, and pandora's problem. In *Conference on Learning Theory, COLT 2021, 15-19 August 2021, Boulder, Colorado, USA*, M. Belkin and S. Kpotufe, Eds. Proceedings of Machine Learning Research, vol. 134. PMLR, 2248–2288.
- HOLMSTROM, B. 1984. On the theory of delegation. *Bayesian Models in Economic Theory*, 115–141.
- HOLMSTROM, B. R. 1978. *On Incentives and Control in Organizations*. Stanford University.
- IMMORLICA, N., LESHNO, J., LO, I., AND LUCIER, B. 2020. Information acquisition in matching markets: The role of price discovery. *Available at SSRN 3705049*.
- KARNI, E. AND SCHWARTZ, A. 1977. Search theory: The case of search with uncertain recall. *Journal of Economic Theory* 16, 1, 38–52.
- KE, T. T. AND VILLAS-BOAS, J. M. 2019. Optimal learning before choice. *Journal of Economic Theory* 180, 383–437.
- KELLER, G. AND OLDALE, A. 2003. Branching bandits: a sequential search process with correlated pay-offs. *J. Econ. Theory* 113, 2, 302–315.
- KLEINBERG, J. M. AND KLEINBERG, R. 2018. Delegated search approximates efficient search. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, É. Tardos, E. Elkind, and R. Vohra, Eds. ACM, 287–302.
- KLEINBERG, R., WAGGONER, B., AND WEYL, E. G. 2016. Descending price optimally coordinates search. In *Proc. 17th ACM Conference on Economics and Computation (EC)*. 23–24. arXiv:1603.07682 [cs.GT].
- LUCIER, B. 2017. An economic view of prophet inequalities. *SIGecom Exch.* 16, 1, 24–47.
- POSTL, P. 2004. Delegated search: Procedure matters. *DEPARTMENT OF ECONOMICS DISCUSSION PAPER-UNIVERSITY OF BIRMINGHAM*.
- SEGEV, D. AND SINGLA, S. 2021. Efficient approximation schemes for stochastic probing and prophet problems. In *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, P. Biró, S. Chawla, and F. Echenique, Eds. ACM, 793–794.
- SINGLA, S. 2018. The price of information in combinatorial optimization. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, A. Czumaj, Ed. SIAM, 2523–2532.
- WEBER, R. 1992. On the gittins index for multiarmed bandits. *The Annals of Applied Probability*, 1024–1033.
- WEISS, G. 1988. Branching bandit processes. *Probability in the Engineering and Informational Sciences* 2, 3, 269–278.
- WEITZMAN, M. L. 1979. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, 641–654.
- WU, C., JIN, C., AND CHEN, Y.-J. 2022. Managing customer search via bundling. *Manufacturing & Service Operations Management* 24, 4, 1906–1925.

Tractable Choice

MODIBO K. CAMARA

University of Chicago

Models in economics and game theory often assume that people behave as if they can solve very complex problems, which can lead to misleading conclusions. To address this, I propose that we supplement the theory of rational choice with a theory of *tractable choice*. Tractable choice asks what an individual can accomplish using resources like time, memory, or data, which are often in short supply. The field of economics has been disciplined when it comes to insisting that choices in models be rational, but is less diligent in requiring that choices be tractable under reasonable assumptions about what resources are available. Fortunately, theoretical computer science has developed deep insights and powerful frameworks for understanding tractability. Using a recent paper as a case study, I argue that tractability is a first-order concern when studying behavior.

Categories and Subject Descriptors: J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms: Economics, Human Factors, Theory

Additional Key Words and Phrases: Bounded Rationality, Computational Complexity

1. INTRODUCTION

Models in economics and game theory often assume that people behave as if they can solve very complex problems. This is concerning if it leads to incorrect predictions about people's behavior. It is also concerning when it comes to designing markets or policies, because the markets and policies that are optimal in our models may be too complicated for real-world actors to interact with. There is a need for a theory that can distinguish predictions and recommendations that are unrealistically complex from those that are at least plausible.

I propose that we supplement the theory of rational choice with a theory of *tractable choice*. To define tractable choice, it is helpful to think of theory as taking a stance on what kinds of behavioral predictions are credible. According to this view, rational choice says that a prediction that “individual i follows strategy s ” is justified only if we can argue that individual i prefers s to any other strategy s' . Tractable choice says that such a prediction is justified only if we can argue that individual i is able to execute strategy s using the resources at her disposal.

Tractable choice asks what an individual can accomplish using resources like time, memory, communication channels, or data, which are often in short supply. Whereas rational choice relies on models of and assumptions about preferences, tractable choice relies on models of these resources and assumptions about their availability. Here, a choice is complex if making said choice requires a large amount of resources. But complexity is multi-faceted. For example, a choice may be complex insofar as the individual must deliberate for a long time, but simple insofar as the individual only needs to communicate a “yes” or “no” answer.

Authors' addresses: mcamara@uchicago.edu

Economics and closely-related fields have been disciplined when it comes to insisting that choices be rational, but are less diligent in requiring that choices be tractable under reasonable assumptions about what resources are available.¹ One reason for this is the lack of consensus on how to model boundedly-rational choice, which asks what people *will* do when faced with intractable problems (as opposed to tractable choice, which simply asks what people *can* do). However, it is not necessary to understand how individuals will respond to intractable problems in order to design markets and policies where optimization is tractable.² Immorlica et al. [2020] is a great example of this approach.

Fortunately, as many readers know well, theoretical computer science has developed deep insights and powerful frameworks for understanding tractability. These frameworks are often highly compatible with economic models, as they tend to be based on a similar foundation of optimization, probability, and logic. They involve general-purpose abstractions that seem capable of representing a wide range of phenomena, not only electronic computers or algorithms implemented in standard programming languages.

It is almost tautological to say that real choices must be tractable, but whether tractability (as understood by computer scientists) should be a first-order concern for the study of human behavior is not quite as obvious. There are three questions that we must ask ourselves:

- (1) Are computational models compatible with and helpful for understanding human behavior?
- (2) Do predictions that respect rationality and tractability look meaningfully different from predictions that only respect rationality?
- (3) Is it really necessary for us to study rational and tractable choice at the same time, rather than having one community (e.g., economists) focus on rationality while another community (e.g., computer scientists) focuses on tractability?

Using my recent paper on “Computationally Tractable Choice” [Camara 2022a] as a case study, I argue that we should expect an affirmative answer to all three questions. My aim is to convince the reader that it is worth taking tractability as seriously as we take rationality, or risk reaching the wrong conclusions. For readers that are already convinced, I hope this case study will help them convince others. In future writing, I hope to address the natural follow-up question of how we can integrate tractability into economic models in a more systematic way.³

¹To a lesser extent, this is also true in algorithmic game theory. For example, there are many models that study auctions or market design from an algorithmic perspective, insisting that allocations can be computed in polynomial time or that the designer’s distributional knowledge come from sample data. But, when it comes to the market participants, many of these models still maintain assumptions like Bayes-Nash equilibrium that are hard to justify as tractable.

²Similarly, it is not necessary to understand precisely how individuals will respond to intractable problems in order to design markets and policies where *approximate* optimization is tractable. We can evaluate such markets and policies according to worst-case participant strategies, subject to the constraint that those strategies be approximately optimal.

³A recent line of work in data-driven mechanism design [Immorlica et al. 2020; Cummings et al. 2020; Camara 2022b; Camara et al. 2020] offers some guidance for modeling tractable choice when data is the limited resource. In addition, a recent line of work by Ryan Oprea develops an

2. COMPUTATIONALLY TRACTABLE CHOICE

Are computational models compatible with and helpful for understanding human behavior? I argue that the answer can be yes, using recent work that integrates computational constraints into decision theory [Camara 2022a]. Still, one must be thoughtful when applying computational models in economics. Developing frameworks that translate results in theoretical computer science to statements relevant for economists seems to require a nuanced understanding of both fields. Echenique et al. [2011] illustrate this point very well.⁴

The premise of Camara [2022a] is that (i) decision-makers have only a limited amount of time to make decisions, but (ii) making good decisions can be time-intensive. To explore the implications for choice, I propose an *axiom of computational tractability*. This axiom is weak: it only rules out behaviors that are thought to be implausible for any algorithm to exhibit in a reasonable amount of time.

I consider a model of choice under risk where the decision-maker has to make many different decisions. For example, consider a consumer choosing from the hundreds or thousands of products in a grocery store, or an investor purchasing shares among the thousands of firms listed on the New York Stock Exchange. The decisionmaker cares about high-dimensional random vectors, i.e.,

$$X = (X_1, \dots, X_n)$$

For example, an investor cares about income X_i from assets $i = 1, \dots, n$, while a consumer cares about consumption bundles, where X_i represents the quantity consumed of good i .

A choice correspondence c maps a menu of feasible options to the decisionmaker's choices X from that menu. The correspondence c is defined over a rich set of menus. This includes all *binary menus* where the decision-maker chooses between two lotteries X and X' , as well as *product menus* where the decision-maker separately chooses each component X_i of the lottery.

I call the choices c *rational* if they maximize expected utility for some utility function u , a common assumption that is axiomatized by von Neumann and Morgenstern [1944]. Later on, I will return to this definition and evaluate its normative appeal in the presence of computational constraints.

I assume that the decisionmaker's choices can be generated by a Turing machine, a powerful model of computation used in computational complexity theory to study what algorithms can and cannot do. Given an appropriate description of a menu, the Turing machine outputs a choice from that menu within a certain amount of time. A choice correspondence is *tractable* if it can be generated by a Turing machine, within an amount of time that grows at most polynomially in the length

experimental paradigm that controls resource complexity while varying incentives, and vice-versa (see e.g., Oprea [2020]). This has already led to some remarkable results, many of which are not yet public, and suggests a path forwards for the empirical study of tractable choice.

⁴Echenique et al. [2011] also integrate computational constraints into decision theory. Their “revealed preference approach to computational complexity” shows that, in a model of consumer choice, any finite and rationalizable dataset can be rationalized by tractable preferences. This surprising result contrasts with the more naive conclusion that consumer choice is intractable because it resembles an NP-hard problem.

of the description.⁵

Having described the model, I can now address two common objections: that humans are not Turing machines, and that computational complexity theory has a misguided focus on worst-case runtime.

The first objection – that humans are not Turing machines – is not a problem in itself. Strictly speaking, it is not necessary for choices to be generated by a Turing machine for the results of Camara [2022a] to hold. All that is necessary is that people are unable to efficiently solve problems that are thought to be fundamentally hard. By contrast, suppose some person *can* make choices that maximize expected utility for a given utility function u , and that those choices are intractable. Then, using the algorithmic reductions developed in the paper, we could leverage that person's choices to efficiently solve problems that are thought to be fundamentally hard. That would be a surprising (and important) result.

The second objection – about worst-case analysis – is best understood as an issue with the definition of rationality, not with the definition of tractability. For context, it is common in computer science to evaluate algorithms by their runtime in the worst-case instance. Consider an algorithm A that takes one minute to solve 99% of inputs and one year for 1% of inputs, so that the worst-case runtime is one year. A decisionmaker that does not have a year to deliberate might use another algorithm A' : see whether A returns an answer within a minute, otherwise choose something suboptimal. This is optimal 99% of the time, suboptimal 1% of the time, and always takes about a minute.

Readers who object to worst-case analysis may point out that the algorithm A' is a perfectly reasonable solution. That may be true. But A' is not rational, insofar as standard definitions of rationality require choice to be optimal 100% of the time (e.g., von Neumann and Morgenstern [1944]). In contrast, A' is tractable because it makes a choice within the time constraint 100% of the time. Moving away from worst-case analysis requires a more flexible definition of rationality, rather than a different definition of tractability.

I use this framework of computationally tractable choice to obtain two kinds of results. First, I show that, under standard rationality assumptions, computational constraints necessarily lead to certain behavioral heuristics. Second, I use these results to give a formal justification for behavior that is not rationalizable by expected utility preferences. I describe these results in the next two sections.

3. FOUNDATIONS FOR BEHAVIORAL HEURISTICS

Do predictions that respect rationality and tractability look meaningfully different from predictions that only respect rationality? In Camara [2022a], I demonstrate that they do look meaningfully different. I show that, under standard rationality assumptions, computational constraints necessarily lead to forms of *choice bracketing*. These are heuristics that lead a decision-maker faced with many decisions

⁵In the paper, I distinguish between weak and strong tractability based on whether the Turing machine has access to polynomial-size advice. I ignore this distinction here and state results informally.

$i = 1, \dots, n$ to focus on each decision i in isolation, without considering the rest.⁶ Equivalently, I show that expected utility maximization is intractable unless the utility function satisfies a strong separability property.

I start by introducing a symmetry assumption that I will later relax. The decisionmaker's choices are *symmetric* if she is indifferent between vectors (X_i, X_j) and (X_j, X_i) . Symmetry may be plausible for investors, where income from one asset i is interchangeable with income from another asset j .

Theorem 1 of Camara [2022a] shows – assuming the $P \neq NP$ conjecture holds – that rational, tractable, and symmetric choices c are observationally equivalent to *narrow choice bracketing*. This means that a decision-maker's choice of X_i in dimension i does not depend on what she chooses in other dimensions j .

More precisely, this result shows that expected utility maximization is intractable unless the utility function is *additively separable*, i.e.

$$u(x) = f(x_1) + \dots + f(x_n)$$

In other words, Theorem 1 is a dichotomy theorem: it partitions a class of computational problems (parameterized by symmetric utility functions u) into polynomial-time (if u is additively separable) and NP-hard (if u is not additively separable).

Theorems 2 and 3 generalize Theorem 1 by dropping the symmetry assumption and strengthening $P \neq NP$ to the *non-uniform exponential time hypothesis*. They show that rational and tractable choice correspondences are observationally equivalent to *dynamic choice bracketing*, a larger class of heuristics that augment choice bracketing with ideas from dynamic programming. These heuristics preserve the computational advantages of choice bracketing while allowing for richer patterns of behavior.

As in Theorem 1, it is useful to restate this characterization in terms of a separability property. Theorem 2 shows that if expected utility maximization is tractable then u is *Hadwiger separable*. This property is a novel relaxation of additive separability that allows for some complementarity and substitutability across dimensions, but limits their frequency. It is quite restrictive and rules out many common utility functions, such as

$$u(x) = f(x_1 + x_2 + \dots)$$

where f is non-linear. More precisely, Hadwiger separability is defined using the notion of an *inseparability graph*. This is an undirected graph where nodes i and j are connected if and only if the utility function u can be represented as

$$u(x_1, x_2, \dots) = f(x_i, x_{-ij}) + g(x_j, x_{-ij})$$

The utility function u is Hadwiger separable if the inseparability graph is sufficiently sparse. That is, if the graph's Hadwiger number grows at most logarithmically in the number of dimensions n .

Together, Theorems 1-3 describe the implications of computational constraints for behavior under standard rationality assumptions. In doing so, they demonstrate that certain behavioral heuristics are not only consistent with but *predicted by* an

⁶There is substantial empirical evidence for this kind of behavior. For example, see Tversky and Kahneman [1981] or Rabin and Weizsäcker [2009] for experimental evidence.

essentially standard model of choice with mild computational constraints. The strength of these results illustrate that tractability can significantly sharpen our predictions about behavior.

Next, we will see that tractability can do more than refine rational choice; it can also highlight problems with how we define rationality in the first place.

4. CHOICE TRILEMMA

Is it really necessary for us to study rational and tractable choice at the same time or is it with minimal loss to have economists focus on rationality and computer scientists focus on tractability? I argue that it is important to study both at the same time. My evidence is the *choice trilemma* of Camara [2022a], which formally shows that incorporating tractability into our models highlights a problem with how we define rationality.

Suppose a decision-maker intrinsically wants to maximize the expected value of a given objective function \bar{u} . If \bar{u} is not Hadwiger separable, Theorem 2 implies that the computationally-constrained decision-maker cannot make choices that exactly optimize the expected value of her objective function. Instead, she might turn to *approximation algorithms* that guarantee her a positive fraction of her optimal payoff. Will this decision-maker make choices that appear rational to an outside observer, insofar as they can be rationalized by some utility function u ?

For many natural objective functions – and assuming $\text{NP} \not\subseteq \text{P/poly}$ – Theorem 4 of Camara [2022a] shows that a computationally-constrained decisionmaker *cannot* simultaneously (i) guarantee any non-zero fraction of her optimal payoff and (ii) be rationalized as maximizing the expected value of some utility function u .

Theorem 4 also shows that the decision-maker can guarantee approximately optimality (i) if she is willing to drop rationality (ii). That is, there do exist tractable algorithms that guarantee at least half of the decision-maker’s optimal payoff. These algorithms do not satisfy the axiomatic definition of rationality of von Neumann and Morgenstern [1944], because they do not exactly maximize the expected value of any particular utility function u .

Altogether, my results imply a *choice trilemma* that relates rationality, tractability, and approximate optimality as properties of choice. For many objective functions \bar{u} , there exist choice correspondences that satisfy any two of these properties, but not all three. That is, a computationally-constrained decision-maker may be better off (according to her true objective function \bar{u}) if she is willing to make choices that an analyst would not be able to rationalize. This suggests that alternative definitions of rationality are needed.

5. CONCLUSION

Using Camara [2022a] as a case study, I argued that tractability should be a first-order concern for economists, and that tools from theoretical computer science can be useful for integrating tractability into economic models.

Specifically, I asked three questions. First, are computational models compatible with and helpful for understanding human behavior? Second, do predictions that respect rationality and tractability look meaningfully different from predictions that only respect rationality? Third, is it really necessary for us to study rational and

tractable choice at the same time? The results in Camara [2022a] suggest that the answers to all three questions are affirmative.

REFERENCES

- CAMARA, M. K. 2022a. Computationally tractable choice. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. EC '22. Association for Computing Machinery, New York, NY, USA, 28.
- CAMARA, M. K. 2022b. Mechanism design with a common dataset. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. EC '22. Association for Computing Machinery, New York, NY, USA, 558.
- CAMARA, M. K., HARTLINE, J. D., AND JOHNSEN, A. 2020. Mechanisms for a no-regret agent: Beyond the common prior. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. 259–270.
- CUMMINGS, R., DEVANUR, N. R., HUANG, Z., AND WANG, X. 2020. Algorithmic price discrimination. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '20.
- ECHENIQUE, F., GOLOVIN, D., AND WIERMAN, A. 2011. A revealed preference approach to computational complexity in economics. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. EC '11. Association for Computing Machinery, New York, NY, USA, 101–110.
- IMMORLICA, N., MAO, J., SLIVKINS, A., AND WU, Z. S. 2020. Incentivizing exploration with selective data disclosure. In *Proceedings of the 21st ACM Conference on Economics and Computation*. EC '20. Association for Computing Machinery, New York, NY, USA, 647–648.
- OPREA, R. 2020. What makes a rule complex? *American Economic Review* 110, 12 (December), 3913–51.
- RABIN, M. AND WEIZSÄCKER, G. 2009. Narrow bracketing and dominated choices. *American Economic Review* 99, 4, 1508–43.
- TVERSKY, A. AND KAHNEMAN, D. 1981. The framing of decisions and the psychology of choice. *Science* 211, 4481, 453–458.
- VON NEUMANN, J. AND MORGENSTERN, O. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.

A Proof of the Nisan-Ronen Conjecture — An Overview

GEORGE CHRISTODOULOU

Aristotle University of Thessaloniki and Archimedes/RC Athena

and

ELIAS KOUTSOUPIS

University of Oxford

and

ANNAMARIA KOVACS

Goethe University, Frankfurt M.

This note presents an overview of our recent publication, which validates a conjecture proposed by Nisan and Ronen in their seminal paper [Nisan and Ronen 2001]. We show that the optimal approximation ratio for deterministic truthful mechanisms for makespan-minimization by a set of n unrelated machines is n .

Categories and Subject Descriptors: F.2 [**Theory of computation**]: Mechanism design

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Algorithmic mechanism design, Nisan-Ronen conjecture

The seminal work [Nisan and Ronen 2001] set the foundations of the field of *algorithmic mechanism design* by probing the computational and information-theoretic limits of *mechanism design*. Mechanism design, a celebrated branch of game theory and microeconomics, studies the design of algorithms (called mechanisms) in environments where the input is privately held and provided by selfish participants. A mechanism for an optimization problem, on top of the traditional algorithmic goal (that assumes knowledge of the input), bears the extra burden of providing incentives to the participants to report their true input. One of the main thrusts in this research area is to demarcate the limitations imposed by truthfulness on algorithms. To what extent are mechanisms less powerful than traditional algorithms?

The objective of the scheduling problem is to minimize the makespan of allocating m tasks to n unrelated machines, where each machine i needs $t_{i,j}$ units of time to process task j . The problem combines various interesting properties. First, it belongs to the most challenging and least explored area of *multi-dimensional* mechanism design, as the private information is multi-dimensional (i.e., for player i , the private values $(t_{i,j})_{j=1}^m$ are a vector). In contrast, the *related* machines scheduling belongs to *single-dimensional* mechanism design, which is well-understood, and for which the power of truthful mechanisms does not substantially differ from the best non-truthful algorithms: not only can truthful mechanisms compute exactly opti-

Authors' addresses:

gichristo@csd.auth.gr, elias@cs.ox.ac.uk, panni@cs.uni-frankfurt.de

mal solutions (if one disregards computational issues [Archer and Tardos 2001]), but a truthful PTAS exists [Christodoulou and Kovács 2013]. Second, the objective of the scheduling problem has a min-max objective, which from the mechanism design point of view is much more challenging than the min-sum objective achievable by the famous VCG [Vickrey 1961; Clarke 1971; Groves 1973] mechanism. VCG is truthful and can be applied to the scheduling problem, but it achieves a very poor approximation ratio, equal to the number of machines n [Nisan and Ronen 2001].

Is there a better mechanism for scheduling than VCG? Nisan and Ronen [Nisan and Ronen 2001] conjectured that the answer should be negative, but for the past two decades, the question has remained open. In our work [Christodoulou et al. 2023] we validate the conjecture.

THEOREM 1. *There is no deterministic truthful mechanism with approximation ratio better than n for the problem of scheduling n unrelated machines.*

Over the years various research attempts with limited success have been made to improve the original lower bound of 2 by Nisan and Ronen. For example, the bound was improved to 2.41 in [Christodoulou et al. 2009], and later to 2.61 in [Koutsoupias and Vidali 2012], which held as the best bound for over a decade. More recently, the lower bound was improved to 2.75 by Giannakopoulos, Hammerl, and Poças [Giannakopoulos et al. 2021], and then to 3 by Dobzinski and Shaulker [Dobzinski and Shaulker 2020]. These improved bounds represented progress in the field, but they left a huge gap between the lower and upper bounds. The first non-constant lower bound for the truthful scheduling problem was given in [Christodoulou et al. 2021a], which showed a lower bound of $\Omega(\sqrt{n})$.

1. THE MAIN ARGUMENT OF THE PROOF

We consider a restricted class of inputs given by a multi-graph where each node is a machine and each edge is a task [Christodoulou et al. 2021b]. For an edge e , we use the notation $e = \{i, j\}$ to denote its vertices i and j , although they do not determine e uniquely. An edge $e = \{i, j\}$ corresponds to a task that has extremely high values for nodes other than i and j , which guarantees that any algorithm with approximation ratio at most n must allocate it to either machine i or machine j (see Figure 1 for an illustration).

The argument deals with multi-cliques with very high *multiplicity*¹, in which *every edge has an endpoint with value 0* (see Figure 3 for an example). The goal is to carefully select a subgraph of this multi-clique and change the values of some of its edges to obtain a lower bound on the approximation ratio. The fact that one of the two values of every edge is 0 is very convenient: a lower bound on the approximation ratio of the subgraph is a lower bound on the approximation ratio of the whole multi-clique as well, since the other edges do not affect the cost of the optimal allocation.

For an instance v , we use the notation v_i^e to denote the value of node i for an edge $e = \{i, j\}$. In most of the argument, we fix the values of the multi-clique and we focus on the boundary functions.

¹The multiplicity of a multi-graph is defined to be the minimum multiplicity among its edges.




Fig. 1. A multi-star instance with three players and four tasks in matrix form (left) and in graph form (right). The symbol ∞ denotes values that are extremely high compared to the other values. This instance is a multi-star, in which player 1 is the root and players 2 and 3 are the leaves.

Definition 2 Boundary function. Fix a mechanism and consider a multi-clique with values v . For an edge $e = \{i, j\}$, the boundary function $\psi_{i,j}^e(z)$ is the threshold value for the allocation of e to node i . More precisely, if we keep all the other values fixed and change the value of e for node j to z , then e is allocated to i if $v_i^e < \psi_{i,j}^e(z)$ and to j if $v_i^e > \psi_{i,j}^e(z)$.

A boundary function $\psi_{i,j}^e(\cdot)$ may depend on the other values of v . Truthfulness severely restricts the class boundary functions. In particular, when we fix all values except the value of a single task, the boundary function $\psi_{i,j}^e(z)$ must be increasing in z . This is single-parameter monotonicity and it is mainly used in the Nice Multi-Star (Theorem 5). A more severe condition on the boundary functions and their relationship comes from multi-parameter truthfulness, that determines how the allocation partitions the space of values (see for example Figure 2). Specifically, the multi-parameter truthfulness for 2 players and 2 tasks plays a central role in the proof and it is repeatedly employed as the main tool for proving the Box Theorem (Theorem 6).

The aim of the proof of the main theorem (Theorem 1) is to show — by the probabilistic method — that there exists a multi-clique of sufficiently high multiplicity that contains a star with approximation ratio at least n , when we keep the values of all other edges fixed. In fact, the argument aims to show that the bound on the approximation ratio for the star is arbitrarily close to $n - 1$. The extra $+1$ in the approximation ratio comes, almost for free, by adding a loop to the root of the star, or equivalently an additional edge between the root and another node j with very high value for j .

To show that there exists a star S with approximation ratio $n - 1$, we roughly aim to show that there exists a star with some root i , with the following properties:

- (1) every edge $e = \{i, j\}$ of S has value 0 for i and the same value z for the leaves, for some $z > 0$.
- (2) the sum of the values of the boundary functions over all edges $\sum_{e \in S} \psi_{i,j}^e(z)$ is at least $(n - 1)z$.
- (3) the mechanism allocates all edges to the root, when we change its values to $\psi_{i,j}^e(z)$ for all $j \neq i$.

It follows immediately that such a star has approximation ratio $n - 1$: the mechanism allocates all tasks to the root with makespan $\sum_{e \in S} \psi_{i,j}^e(z) \geq (n - 1)z$, while a better allocation is to allocate all tasks to the leaves with makespan z .

A star that satisfies the second property will be called *nice* and the third property *box*. In [Christodoulou et al. 2021a], we used an argument that is similar to the

Box Theorem, which establishes the box property for many stars. Actually the Box Theorem is a cleaner and stronger argument than the one in [Christodoulou et al. 2021a], and one can use it to obtain the results of that work directly. To completely resolve the Nisan-Ronen conjecture we needed to work with multi-cliques, not only stars as in [Christodoulou et al. 2021a], and the Nice Multi-Star Theorem allows us to focus on a particular multi-star of the multi-clique.

For technical reasons we need to work with approximate notions of niceness and box-ness.

Definition 3 Nice star and nice multi-star. For a given $\varepsilon > 0$ and an instance v , a star S with root i and leaves all the remaining $(n - 1)$ nodes is called ε -nice, or simply nice, if there exists $z > 0$ such that:

- (1) every edge $e = \{i, j\}$ of S has value $v_i^e = 0$ for root i and $v_j^e \in (z, (1 + \varepsilon)z)$ for leaf j
- (2)

$$\sum_{e \in S} \psi_{i,j}^e(v_j^e) \geq (1 - 3\varepsilon)(n - 1)z. \quad (1)$$

A multi-star is nice if *all of its stars* with $n - 1$ leaves are nice, with the *same* z .

By letting ε tend to 0, v_j^e can be arbitrarily close to z . Next we define boxes².

Definition 4 Box. For a given $\delta > 0$ and instance v , a star S with root i is called δ -box, or simply box, if every edge $e = \{i, j\}$ of S has value 0 for i and the mechanism allocates all edges to i , when we change their value for i to $\psi_{i,j}^e(v_j^e) - \delta$ for every leaf j of S .

Now that we have the definitions of nice multi-stars and box stars, we can state the two main theorems that almost immediately establish the main result. The first theorem establishes the existence of nice multi-stars of arbitrarily high multiplicity (see Figure 3). The second theorem asserts that nice multi-stars with sufficiently high multiplicity contain a box star of $n - 1$ leaves (see Figure 3).

THEOREM 5 NICE MULTI-STAR. *For every mechanism with bounded approximation ratio and every q , there exists a multi-clique that contains a nice multi-star with multiplicity q .*

THEOREM 6 BOX. *Fix $\delta, \varepsilon > 0$ and a mechanism with approximation ratio at most n . Consider an instance that contains a multi-star, of sufficiently high multiplicity, in which all values of the root i are 0 and all values of the leaves are in $(z, (1 + \varepsilon)z)$. Then the multi-star contains a star with $n - 1$ leaves, which is a δ -box.*

The proof of the main result (Theorem 1) follows immediately from the above two theorems. Use Theorem 5 to find a multi-clique that contains a nice multi-star with sufficiently high multiplicity. Use Theorem 6 to find a nice box inside it. The next lemma makes this precise.

²See Figure 2 for an illustration.

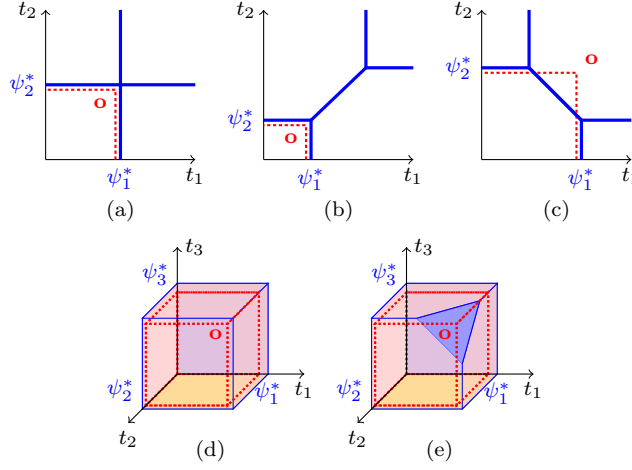


Fig. 2. (Box). Allocation partitions of the root values for a star of 2 leaves (a)-(c) and 3 leaves (d)-(e); in the latter case, only part of the allocation partition is shown. Call the root i and leaves $j \in \{1, 2, 3\}$. If we denote the edges of the star by e_j , the figure uses the shorthand: $t_j = v_i^{e_j}$ for the values of the root, and $\psi_j^* = \psi_{i,j}^{e_j}(v_j^{e_j})$ for the boundary values. Truthfulness restricts the shapes and boundaries of the allocation areas. The dotted red lines correspond to values $\psi_j^* - \delta$ of the box definition. Cases (a), (b), and (d) are boxes, as the corner o is inside the region where the root gets all the tasks. On the other hand, cases (c) and (e) are not boxes, since the corner point o lies outside this region.




Fig. 3. This is an illustration of the components that appear in the statement of Theorem 1. (a) shows a multi-clique with $n = 4$ nodes and multiplicity 6. It should be noted that the actual multiplicity needed is much higher. For simplification, we use z to denote non-zero values, which are not necessarily the same for all edges. (b) shows a multi-star which is a subgraph of the multi-clique. If this is a *nice* multi-star, the value z is approximately the same for all leaves. (c) shows a simple star of this nice multi-star, selected by the Box Theorem 6. The nice-ness property roughly guarantees that $\psi_1 + \psi_2 + \psi_3 \geq 3z$, and the box-ness property guarantees that all tasks will be allocated to the root. This will give approximation ratio roughly 3; we can increase this to $n = 4$ by adding loops. Note that the remaining edges — which do not appear in (c) — do not contribute to the optimal makespan, because one of their values is 0.

LEMMA 7. *A δ -nice box with a loop in the root, in which all values of the root are 0 and all values of the leaves are in $(z, (1 + \varepsilon)z)$, has approximation ratio n , as δ and ε tend to 0.*

PROOF. Take a nice box and consider the instance when we change the values of the root i to $\psi_{i,j}^e(v_j^e) - \delta$ for all $j \neq i$. By the box-ness property all the edges are allocated to the root. Change now the value of the loop to z and decrease the values of the root to $\psi_{i,j}^e(v_j^e) - 2\delta$. The task that corresponds to the loop must still be allocated to root, even when we increase its value to z . By applying monotonicity, the allocation of the edges remains the same. The makespan of the mechanism is

$$z + \sum_{j \neq i} (\psi_{i,j}^e(v_j^e) - 2\delta) \geq z + (1 - 3\varepsilon)(n - 1)z - 2(n - 1)\delta,$$

while the optimal makespan is at most $(1 + \varepsilon)z$, (when the root gets the loop and the leaves get the remaining edges). The ratio tends to n as δ and ε tend to 0. \square

2. CONCLUSION

Our work [Christodoulou et al. 2023] validates the Nisan-Ronen conjecture, by establishing a lower bound for all deterministic truthful mechanisms. The case of randomized or fractional mechanisms is still open and it appears to be challenging; the best known lower bound of the approximation ratio is 2 [Mu’alem and Schapira 2018; Christodoulou et al. 2010], significantly lower than the best known upper bound $(n + 1)/2$. The bottleneck of applying the techniques of the current work to these variants appears to be the lack of a good characterization of 2×2 fractional mechanisms. Another important direction is to apply our approach to major open questions in other settings and in particular to combinatorial auctions.

REFERENCES

- ARCHER, A. AND TARDOS, É. 2001. Truthful mechanisms for one-parameter agents. In *FOCS*. IEEE Computer Society, 482–491.
- CHRISTODOULOU, G., KOUTSOPIAS, E., AND KOVÁCS, A. 2010. Mechanism design for fractional scheduling on unrelated machines. *ACM Transactions on Algorithms* 6, 2.
- CHRISTODOULOU, G., KOUTSOPIAS, E., AND KOVÁCS, A. 2021a. On the Nisan-Ronen conjecture. In *FOCS*. IEEE, 839–850.
- CHRISTODOULOU, G., KOUTSOPIAS, E., AND KOVÁCS, A. 2021b. Truthful allocation in graphs and hypergraphs. In *ICALP*. LIPIcs, vol. 198. 56:1–56:20.
- CHRISTODOULOU, G., KOUTSOPIAS, E., AND KOVÁCS, A. 2023. A proof of the Nisan-Ronen conjecture. In *STOC*. ACM, 672–685.
- CHRISTODOULOU, G., KOUTSOPIAS, E., AND VIDALI, A. 2009. A lower bound for scheduling mechanisms. *Algorithmica* 55, 4, 729–740.
- CHRISTODOULOU, G. AND KOVÁCS, A. 2013. A deterministic truthful ptas for scheduling related machines. *SIAM J. Comput.* 42, 4, 1572–1595.
- CLARKE, E. H. 1971. Multipart pricing of public goods. *Public Choice* 8.
- DOBZINSKI, S. AND SHAULKER, A. 2020. Improved Lower Bounds for Truthful Scheduling.
- GIANNAKOPOULOS, Y., HAMMERL, A., AND POÇAS, D. 2021. A new lower bound for deterministic truthful scheduling. *Algorithmica* 83, 9, 2895–2913.
- GROVES, T. 1973. Incentives in teams. *Econometrica* 41, 4, 617–631.
- KOUTSOPIAS, E. AND VIDALI, A. 2012. A lower bound of $1 + \phi$ for truthful scheduling mechanisms. *Algorithmica*, 1–13.

- MU'ALEM, A. AND SCHAPIRA, M. 2018. Setting lower bounds on truthfulness. *Games and Economic Behavior* 110, 174–193.
- NISAN, N. AND RONEN, A. 2001. Algorithmic mechanism design. *Games and Economic Behavior* 35, 166–196.
- VICKREY, W. 1961. Counterspeculations, auctions and competitive sealed tenders. *Journal of Finance* 16, 8–37.

Deep Reinforcement Learning for Economics: Progress and Challenges

ETAN A. GREEN

University of Pennsylvania, Arena–AI
and

E. BARRY PLUNKETT

Skip Protocol

We discuss the application of deep reinforcement learning to economic domains in general, and to bargaining on eBay in particular.

Categories and Subject Descriptors: J.4 [Social and Behavioral Science]: Economics

General Terms: Economics, Algorithms

Additional Key Words and Phrases: Deep reinforcement learning, Bargaining, Offline RL

1. INTRODUCTION

If, in 2018, you had asked about the most promising advance in artificial intelligence, the answer almost certainly would have been deep reinforcement learning. AlphaZero, trained using deep RL, had just been crowned the world’s best player of chess, Go, and Shogi [Silver et al. 2018], and the application to real-world domains seemed imminent. “Artificial intelligence,” said David Silver, winner of the 2019 ACM Prize in Computing, “is deep reinforcement learning.” [Silver 2016]

Today, the promise of deep RL has not been realized. The fundamental challenge is that reinforcement learning agents require an environment in which to train, and creating an environment that reliably simulates the real world has proven difficult. The stories of RL successes are almost universally stories of pre-existing, reliable training grounds. This letter discusses an exception—an RL agent that bargains on eBay [Green and Plunkett 2022]—and the promise, as well as the challenges, it portends for applications of RL in economic domains and in the real world more generally. Our view is optimistic, if somewhat dystopian: in the near future, many economic decisions will be made by reinforcement learning agents.

2. BACKGROUND

Reinforcement learning agents learn by trial and error. They observe the state of the world (e.g., the board position in chess), take an action (a move) in a given state (board position), receive a reward (based on the outcome of the game), and reinforce actions that lead to higher rewards. To learn, RL agents need an environment that communicates the consequence of an action: the state in which it will take an action next, and the reward it receives for arriving at that state. By traversing many—often millions, sometimes billions—of states, the agent can learn a policy:

Authors’ addresses: etangreen@gmail.com, bpiv400@gmail.com

an action to take in every state that maximizes a potentially distant reward, e.g., a move to make in every board position that maximizes the probability of winning. In deep RL, the mapping from states to actions is learned by a neural network.

The most ready domains for RL are those in which a reliable environment already exists. Among the most popular settings for testing RL algorithms are Atari games, for which the environment is the game itself [Hafner et al. 2019]. The state is defined by the pixels on the screen, an action is a move of the joystick, and the reward is the score. When the agent takes an action, the game responds by updating the pixels on the screen and the player’s score.

Recent RL-driven advances in algorithms for matrix multiplication [Fawzi et al. 2022] and sorting [Mankowitz et al. 2023] also exploit pre-existing environments. For sorting, the state is the current order of elements in the array, an action may swap two elements, for instance, and the reward is a penalty for each action taken (so that the optimal policy is one that sorts an array in the fewest number of actions). When the agent takes an action, the next state is simply the new ordering of the array.

Adversarial games like chess pose an added complication: after the agent acts, the state in which it acts next depends on how the opponent responds. Hence, the training environment must incorporate the opponent’s response. This problem neatly disappears in two-player, zero-sum games, such as chess and Go. To act optimally in these games, an agent need not learn to best respond to any opponent. Rather, it is sufficient to learn a best response to a best-responding opponent, which an RL agent can learn by playing against itself. By virtue of the minimax theorem, an equilibrium strategy learned in this manner will be optimal against any opponent, regardless of their intelligence. In a matter of days, AlphaZero went from knowing nothing about chess save the rules to the best chess player in the world simply by playing against itself millions of times.

3. BARGAINING ON EBAY

Many real-world games, and particularly those of economic interest, are neither two-player nor zero-sum. Bargaining, for instance, is multi-player: a seller may bargain with more than one buyer. It is also not zero-sum: the buyer and seller share a surplus only if they reach an agreement; otherwise, no surplus is generated. In multi-player or non-zero sum games, the goodness of an action depends on the opponent. Policies that perform well against one opponent may perform poorly against another.

One way around this problem is to train agents that perform well against a particular type of opponent: humans. An agent that exploits humans is useful in two ways: first, to exploit humans in the real world; and second, to help humans make better decisions.

We trained a deep RL agent to exploit humans when bargaining on eBay (in Best Offer listings, in which a seller sets a list price, and buyers and sellers may negotiate a lower price by making offers sequentially) [Green and Plunkett 2022]. The strategy that the agent learned, as either the buyer or the seller, meaningfully outperforms those that humans play. As the seller, the agent sells items more often and for higher prices. As the buyer, the agent purchases items more often and for

lower prices.

We show that most of these gains can be attained through simple tactics. For instance, the seller exploits human buyers by rejecting most first offers, particularly generous first offers, or those that request only a small discount on the list price. Generous first offers signal the buyer's willingness to pay more. By rejecting such offers, the seller communicates that the list price is firm. Human buyers often respond by paying the full list price.

The primary methodological contribution of our work is a template for training RL agents to exploit humans in real-world economic games. In a perfect world, we would have trained the agent on eBay—i.e., by making offers, observing counteroffers, and reinforcing offers that lead to higher payoffs. However, deep RL algorithms require an impossibly large number of actions to learn intelligent policies. We could neither list millions of items on eBay nor make millions of offers.

Our solution was to train a model of the real world from a massive dataset of negotiations on eBay [Backus et al. 2020], and then to train an RL agent in that model. This environment model is a neural network that simulates human behavior on eBay. The model predicts when buyers arrive, what offers they make, and how sellers respond—conditional on the features of the listing and the sequence of prior offers. When the agent acts as a buyer, we sample seller counteroffers from the environment. When the agent acts as a seller, we sample buyer arrivals, first offers, and counteroffers. In this manner, the agent bargains against millions of (simulated) humans in a couple days, and for only the cost of that compute time.

This approach is not without challenges, the most fundamental of which is that the environment model may not perfectly correspond to the real world. This difficulty has impeded applications of RL in robotics, in which agents are often trained in a model of the physical environment. Some aspects of the physical world, such as friction and wear on robotic arms, are difficult to model. As a result, agents trained to perform tasks in the model often cannot perform those tasks in the world [Kormushev et al. 2013].

4. CHALLENGES

Economic domains pose an added difficulty: confoundedness, or missing data. For instance, a buyer or seller on eBay may attach a text message to their offer, and while the eBay dataset we use contains an indicator for whether a text message accompanied the offer, it does not contain the content of the message. What someone says in their message probably affects how the other party responds to their offer [Backus et al. 2021]. If, say, nice messages induce acceptances and mean ones induce rejections, our model will sometimes respond with an acceptance and other times with a rejection—not because the true distribution is bimodal but because messages are sometimes nice and sometimes mean.

A second challenge concerns exploration. Often during training, an RL agent will try an unusual action, such as offering \$0. In the real world, a seller will learn that an offer of \$0 is unprofitable, and a buyer will learn that it is a waste of time. However, offers of \$0 do not exist in our training data because humans never make them; hence, there is no guarantee that the model we train from those data will learn these truths.

We circumvent this issue by restricting the offers that the RL agent can make to those that are common in the data. Because the game tree is shallow—eBay limits the buyer and seller to no more than three offers each—this constraint mostly keeps the agent within the distribution of the training data. In settings with deeper game trees, however, exploration will lead to novel states, even if the training data are large and varied.

A more sophisticated approach is to penalize the agent for exploring outside the confines of the training data. This can be done by penalizing actions that the environment model deems unlikely, or by training an ensemble of environment models, each on a different partition of the training data, and penalizing actions that induce disagreement among the models. Neither approach seems to work well in practice. Rather than smoothly converging to a policy that balances rewards and penalties, standard RL algorithms like PPO oscillate between maximizing reward and minimizing exploration penalty without ever converging [Moskovitz et al. 2023].

To this point, we have considered RL approaches that learn *online*, by training either in the environment of interest or by training in a model of that environment. A newer, more promising alternative may be offline RL, in which a policy is learned directly from data [Kostrikov et al. 2021]. Offline RL proceeds in two steps. First, the data are used to train a critic, or a neural network that estimates the sum of discounted future rewards for taking a given action in the current state, and then taking the best sequence of actions observed in the data. Second, a policy is extracted from the critic by finding the best action in each state. One way to do this is to first train a *clone*, or a model that predicts the distribution of actions taken in the data, conditional on the state. Sampling actions from the clone yields a set of candidate actions. Evaluating those actions using the critic identifies which is best. This approach more naturally constrains exploration to actions that are in the distribution of the training data.

By its very name, offline RL offers an alternative to online RL. However, we view these approaches as complementary. Since the offline critic and the environment model both process state-action pairs, they can share a neural architecture. Hence, they can be trained jointly, by adding their losses before backpropagation. This conjoined approach may yield a better critic—by forcing the model to predict state transitions explicitly, rather than simply their rewards. A second advantage of training an environment model alongside an offline critic is that the environment model can be used to evaluate the policy extracted from the critic.

5. CONCLUSION

The challenges of applying deep reinforcement learning to real-world economic problems are significant, but so are the rewards. Methodological advancements offer hope that this promise will soon be realized.

REFERENCES

- BACKUS, M., BLAKE, T., LARSEN, B., AND TADELIS, S. 2020. Sequential bargaining in the field: Evidence from millions of online bargaining interactions. *The Quarterly Journal of Economics* 135, 3, 1319–1361.
- BACKUS, M., BLAKE, T., PETTUS, J., AND TADELIS, S. 2021. Communication and bargaining breakdown: An empirical analysis. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 129–129.

- FAWZI, A., BALOG, M., HUANG, A., HUBERT, T., ROMERA-PAREDES, B., BAREKATAIN, M., NOVIKOV, A., R. RUIZ, F. J., SCHRITTWIESER, J., SWIRSZCZ, G., SILVER, D., HASSABIS, D., AND KOHLI, P. 2022. Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature* 610, 7930, 47–53.
- GREEN, E. A. AND PLUNKETT, E. B. 2022. The science of the deal: Optimal bargaining on ebay using deep reinforcement learning. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 1–27.
- HAFNER, D., LILLICRAP, T., BA, J., AND NOROUZI, M. 2019. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*.
- KORMUSHEV, P., CALINON, S., AND CALDWELL, D. G. 2013. Reinforcement learning in robotics: Applications and real-world challenges. *Robotics* 2, 3, 122–148.
- KOSTRIKOV, I., NAIR, A., AND LEVINE, S. 2021. Offline reinforcement learning with implicit q-learning. *arXiv preprint arXiv:2110.06169*.
- MANKOWITZ, D. J., MICH, A., ZHERNOV, A., GELMI, M., SELVI, M., PADURARU, C., LEURENT, E., IQBAL, S., LESPIAU, J.-B., AHERN, A., KÖPPE, T., MILLIKIN, K., GAFFNEY, S., ELSTER, S., BROSHEAR, J., GAMBLE, C., MILAN, K., TUNG, R., HWANG, M., CEMGIL, T., BAREKATAIN, M., LI, Y., MANDHANE, A., HUBERT, T., SCHRITTWIESER, J., HASSABIS, D., KOHLI, P., RIED-MILLER, M., VINYALS, O., AND SILVER, D. 2023. Faster sorting algorithms discovered using deep reinforcement learning. *Nature* 618, 7964, 257–263.
- MOSKOVITZ, T., O'DONOGHUE, B., VEERIAH, V., FLENNERHAG, S., SINGH, S., AND ZAHAVY, T. 2023. Reload: Reinforcement learning with optimistic ascent-descent for last-iterate convergence in constrained mdps. *arXiv preprint arXiv:2302.01275*.
- SILVER, D. 2016. Tutorial: Deep reinforcement learning. https://icml.cc/2016/tutorials/deep_rl_tutorial.pdf. Accessed: 2023–06-09.
- SILVER, D., HUBERT, T., SCHRITTWIESER, J., ANTONOGLOU, I., LAI, M., GUEZ, A., LANCTOT, M., SIFRE, L., KUMARAN, D., GRAEPEL, T., ET AL. 2018. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science* 362, 6419, 1140–1144.

Mechanism Design with Predictions: An Annotated Reading List

ERIC BALKANSKI

Columbia University

and

VASILIS GKATZELIS

Drexel University

and

XIZHI TAN

Drexel University

A surge of recent work has focused on analyzing the performance of algorithms guided by predictions, aiming to enhance their worst-case performance guarantees with improved guarantees when the predictions are accurate. This “learning-augmented” framework was recently also extended to mechanism design settings involving strategic agents and we provide an overview of these results.

Categories and Subject Descriptors: B.6.3 [**Theory of computing**]: Algorithmic mechanism design

General Terms: Mechanism Design, Beyond Worst-Case Analysis

Additional Key Words and Phrases: Consistency, Robustness

For more than half a century, the dominant approach for the mathematical analysis of algorithms in computer science has been worst-case analysis. While worst-case analysis provides a useful signal regarding the robustness of an algorithm, it can be overly pessimistic and it often leads to uninformative bounds or impossibility results that may not reflect real-world obstacles. Meanwhile, advances in machine learning have led to very practical algorithms, most of which do not provide any non-trivial worst-case performance guarantees. Motivated by the tension between worst-case analysis and machine learning, a surge of recent work aims to design robust algorithms guided by machine-learned predictions. The goal of this literature on “algorithms with predictions” is to *simultaneously* provide two types of guarantees: “robustness” (which corresponds to the classic worst-case guarantees, even if the predictions are arbitrarily bad) and “consistency” (i.e., the performance guarantees when the predictions are accurate). This “learning-augmented framework” has been used successfully in a variety of settings, e.g., toward a refined analysis of competitive ratios in online algorithms and running times in traditional algorithms.

A very recent line of work has deployed this learning augmented framework in settings involving strategic agents. In such settings, the designer often faces information limitations, e.g., the participating agents’ may have private information which they can strategically misreport, which limits the designer’s ability to reach

Authors’ addresses: eb3224@columbia.edu, gkatz@drexel.edu, xizhi@drexel.edu

desired outcomes. Mechanism design has proposed solutions to this problem, but their worst-case guarantees are often underwhelming from a practical perspective. Could we design learning-augmented mechanisms that combine “robustness” with strong “consistency” guarantees? Below are some initial works in this direction.

- (1) Priyank Agrawal, Eric Balkanski, Vasilis Gkatzelis, and Xizhi Tan. “Learning-Augmented Mechanism Design: Leveraging Predictions for Facility Location”. In: *EC ’22: The 23rd ACM Conference on Economics and Computation*. ACM, 2022, pp. 497–528

This work initiated the line of research on learning-augmented mechanism design and showcased the power of predictions in strategic settings, focusing on the canonical problem of strategic facility location. For both the egalitarian social cost and utilitarian cost, this paper provided truthful mechanisms enhanced with a prediction regarding the optimal facility location. These mechanisms achieve either the optimal consistency with the best-possible robustness (i.e., the best of both worlds) or the optimal trade-off between the two notions.

- (2) Chenyang Xu and Pinyan Lu. “Mechanism Design with Predictions”. In: *IJCAI ’22: The 31st International Joint Conference on Artificial Intelligence*, 2022, pp. 571–577

This work concurrently initiated mechanism design with predictions. Rather than focusing on a single problem, it sampled a variety of different mechanism design problems, including auction design, frugality, scheduling, and facility location. In each of these settings, the results are truthful mechanisms that utilize predictions to achieve consistency guarantees that are better than the best-known worst-case performance guarantees, while simultaneously maintaining non-trivial robustness guarantees.

- (3) Vasilis Gkatzelis, Kostas Kollias, Alkmini Sgouritsa, and Xizhi Tan. “Improved Price of Anarchy via Predictions”. In: *EC ’22: The 23rd ACM Conference on Economics and Computation*. ACM, 2022, pp. 529–557

While most papers on learning-augmented mechanism design focus on centralized mechanisms, this work studies a decentralized setting where the mechanism has limited information and can affect the agents’ decisions only indirectly. It proposes cost-sharing protocols for classic job scheduling and network creation games which use predictions regarding the missing information and induce better Nash equilibria and improved price of anarchy bounds.

- (4) Eric Balkanski, Vasilis Gkatzelis, and Xizhi Tan. “Strategyproof Scheduling with Predictions”. In: *ITCS ’23: The 14th Innovations in Theoretical Computer Science Conference*, vol. 251. 2023, 11:1–11:22

This work focused on the celebrated problem of makespan minimization in strategic scheduling introduced by one of the first papers in AGT. It was con-

jectured, and very recently validated, that the best deterministic mechanism cannot achieve an approximation better than n . In this work, the authors provided a polynomial time mechanism, enhanced with predictions, that is 6-consistent and $2n$ -robust, thus achieving asymptotically the best of both worlds (asymptotically optimal consistency and robustness).

- (5) Gabriel Istrate and Cosmin Bonchis. “Mechanism Design With Predictions for Obnoxious Facility Location”. In: *CoRR* abs/2212.09521 (2022)

This work considers the obnoxious facility location problem where the agents wish to be far from the facility instead of close to it. For segments, squares, circles, and trees, the authors provide truthful mechanisms augmented with predictions and bounded their consistency and robustness. The trade-offs obtained are shown to be optimal in one dimension.

- (6) Maria-Florina Balcan, Siddharth Prasad, and Tuomas Sandholm. “Bicriteria Multidimensional Mechanism Design with Side Information”. In: *CoRR* abs/2302.14234 (2023)

This work focuses on multidimensional mechanism design. Rather than focusing on any specific setting, it proposes a general meta-mechanism that incorporates different types of side information to achieve both high social welfare and high revenue. The approach is versatile and can accommodate various sources of side information.

- (7) Andres Muñoz Medina and Sergei Vassilvitskii. “Revenue Optimization with Approximate Bid Predictions”. In: *NIPS '17: Advances in Neural Information Processing Systems 30*. 2017, pp. 1858–1866

This closely related work appeared before the papers initiating the line of work on algorithms with predictions. It focuses on finding good reserve prices in advertising auctions and proposes a method to reduce reserve price optimization to a standard setting of prediction under squared loss. They used a predictor to define a clustering of the data and compute the empirically maximizing reserve price for each group. The reduction directly ties the revenue gained by the algorithm to the prediction error, but without bounded robustness guarantees.

- (8) Michael Mitzenmacher and Sergei Vassilvitskii. “Algorithms with Predictions”. In: *Commun. ACM* 65.7 (2022), pp. 33–35

A survey of some of the initial results on algorithms with predictions.

- (9) Alexander Lindermayr and Nicole Megow. *Algorithms with Predictions*. URL: <https://algorithms-with-predictions.github.io/>

This frequently updated website keeps track and categorizes papers in the area of algorithms with predictions. It allows easy search of papers by performance measure and/or type of problem.

Solution to Exchanges 20.1 Puzzle: Communicating to Plan Noam Nisan’s 60th Birthday Workshop

ALEC SUN

Carnegie Mellon University

This is a solution to Vincent Conitzer’s puzzle “Communicating to Plan Noam Nisan’s 60th Birthday Workshop”, which appeared as below in the July 2022 issue of SIGecom Exchanges.

Michael, Moshe, and Shahar—i.e., a constant number of organizers—are planning the workshop for Noam’s 60th birthday, and are trying to predict who, out of n people, will attend. Whether a person wants to attend is a function of who else attends. “The more the merrier,” so for each person i , if i would attend when S is the set of other attendees, and $S \subseteq S'$, then i would attend when S' is the set of other attendees. Let \mathcal{S}_i be the set of sets S of other people for which i would attend (so, \mathcal{S}_i is upward closed).

To split the work, the organizers partition the set of n people among themselves. Subsequently, each of them figures out, for every player i in his own part, what \mathcal{S}_i is. (Note that each organizer thus still needs to think about how much “his” people like the people in the other parts. But each organizer knows \mathcal{S}_i only for people i in his own part.) At this point, the organizers, who of course want the workshop to be successful, must communicate with each other to find the *largest* possible set of people S^* that can consistently attend (i.e., the largest set with the property such that every person in it will attend given that everyone else in the set attends: i.e., for each $i \in S^*$, we have $S^* \setminus \{i\} \in \mathcal{S}_i$, and S^* is the largest set with that property).

Up to a constant factor, how many bits of communication do the organizers need to figure this out?

We claim that the organizers need $\Theta(n \log n)$ bits of communication.

1. UPPER BOUND

We begin by asking the following question. When would someone not attend the birthday celebration? Certainly if $\mathcal{S}_i = \emptyset$ then person i would not attend. It turns out that this is the only possible restriction preventing someone from attending.

CLAIM 1.1. *If $\mathcal{S}_i \neq \emptyset$ for all i , that is, for every person i there is at least one set S of other attendees for which i would attend, then $S^* = [n]$.*

Claim 1.1 follows directly from “the more the merrier” and implies the following recursive algorithm for the birthday problem:

- If an organizer sees that $\mathcal{S}_i = \emptyset$ for a person i in their part, they communicate the identity of i to the other organizers using $O(\log n)$ bits and then remove person i from consideration.
- Since person i provably cannot attend, all of the organizers remove all sets containing i from each \mathcal{S}_j in their parts. Note that the \mathcal{S}_j remain upward-closed and hence we have reduced to an instance of the same problem with $n - 1$ people.
- Repeat the above procedure until among the people currently in consideration, which is possibly the empty set, there is no person i for which $\mathcal{S}_i \neq \emptyset$.

Authors’ addresses: sundogx@gmail.com

Claim 1.1 shows that S^* is precisely the set of people remaining after the above algorithm is run. Since the identity of each person is communicated at most once in the above procedure, the total communication is $O(n \log n)$.

We remark that the above procedure is reminiscent of Moulin’s mechanism for cost-sharing in which a designer decides on which players to be served and what cost to charge them through an iterative process. As long as there exists at least one player that has a cost-share strictly greater than their bid, that player is removed from consideration and new cost-shares are computed with the remaining players. This process is continued until all remaining players’ cost-shares are at most their bids, at which point the mechanism terminates and serves this remaining set of players. See [Moulin 1999] and [Moulin and Shenker 2001] for more details on Moulin’s mechanism.

2. LOWER BOUND

Now we prove that $\Omega(n \log n)$ bits of communication are required to compute S^* . For simplicity, suppose there are only two organizers, Michael and Moshe, and that the number n of people is even. Number the people $1, 2, \dots, n$ and suppose that Michael knows S_i for all odd i , which we denote by Michael’s input x to the problem, and Moshe knows S_i for all even i , which we denote by Moshe’s input y . We call the pair (x, y) , which is the aggregated set of preferences, the *input* to the problem, and we call the resulting set S^* the *answer*.

Consider the *communication transcript* of the organizers, which consists of all bits communicated between them as well as the final answer. Without loss of generality, assume that the two players alternate in communicating bits, with Michael communicating first. We want to show that the communication transcript must have size $\Omega(n \log n)$. We use the *fooling set* method, a lower bound technique in communication complexity that appears in Nisan’s own book, co-authored with Kushilevitz, on the subject [Kushilevitz and Nisan 1996]. The idea of the fooling set method is that if two distinct input pairs have the same communication transcript, then we can find two other pairs that also have this same transcript.

CLAIM 2.1. *Let (x, y) and (x', y') be two inputs to Michael and Moshe that have the same communication transcript. Then the inputs (x', y) and (x, y') also have this same communication transcript.*

PROOF. Each bit communicated by a player is a deterministic function of that player’s input and the bits seen so far. Since (x, y) and (x', y') have the same transcript, the first bit b_1 sent, which depends only on Michael’s input, is the same whether Michael has input x or x' , so the first bits of the transcripts of (x', y) and (x, y') are also b_1 . The second bit b_2 sent is a function of b_1 , which we already argued is the same in all four transcripts, and Moshe’s input. Since the transcripts of (x, y) and (x', y') have the same second bit b_2 , (x', y) and (x, y') also have b_2 as their second bits. Continuing inductively, we deduce that (x, y) , (x', y') , (x', y) , (x, y') all produce identical transcripts. \square

Our strategy will be to find many input pairs $(x_1, y_1), (x_2, y_2), \dots, (x_m, y_m)$ all with the same answer S^* . By the Pigeonhole Principle, if m is large enough and the amount of communication is limited, then two of these input pairs $(x_i, y_i), (x_j, y_j)$ will have the same communication transcripts. By Claim 2.1, the input pairs (x_i, y_j) and (x_j, y_i) will also have this same communication transcript and in particular the same answer. However, if we had constructed these inputs from the start in such a way that all inputs (x_i, y_j) for $i \neq j$ actually had answers that are different than S^* , then this would yield a contradiction.

Consider all *parity-alternating* permutations $\sigma = \sigma_1 \sigma_2 \dots \sigma_n$ of the n people such that σ_1 is odd, σ_2 is even, and so on. Note that there are $\frac{n}{2}$ ways to choose σ_1 since there are

$\frac{n}{2}$ odd people, $\frac{n}{2}$ ways to choose σ_2 since there are $\frac{n}{2}$ even people, $\frac{n}{2} - 1$ ways to choose σ_3 since there are $\frac{n}{2} - 1$ remaining odd people, and so on, for a total of $\left(\left(\frac{n}{2}\right)!\right)^2$ such permutations.

We say that each parity-alternating permutation $\sigma = \sigma_1 \dots \sigma_n$ induces an input to the birthday problem as follows. Let $\mathcal{S}_{\sigma_i} = \{S : \exists j < i, \sigma_j \in S\}$, that is, the preference of person σ_i is “I would attend only if at least one of $\sigma_1, \dots, \sigma_{i-1}$ attends.” For $i = 1$ this simply means that person σ_1 would not attend. We note that the \mathcal{S}_{σ_i} are indeed upward-closed. For all inputs (x, y) induced by such permutations σ , we have $S^* = \emptyset$ since σ_1 would not attend, which prohibits σ_2 from attending, and so on.

CLAIM 2.2. *For any two distinct parity-alternating permutations σ and σ' , which induce preferences (x, y) and (x', y') respectively as inputs to Michael and Moshe, the answers to the inputs (x', y) and (x, y') are not $S^* = \emptyset$.*

EXAMPLE 2.3. *Consider the permutations $\sigma = 1234$ and $\sigma' = 1432$, which induce the inputs*

$x =$ “1 would not attend”, “3 would only attend if 1 or 2 does”
 $y =$ “2 would attend only if 1 does”, “4 would attend only if 1 or 2 or 3 does”
 $x' =$ “1 would not attend”, “3 would attend only if 1 or 4 does”
 $y' =$ “2 would attend only if 1 or 4 or 3 does”, “4 would attend only if 1 does”.

Note that if the input was (x', y) then 3 and 4 can attend together, and if the input was (x, y') then 2 and 3 can attend together.

PROOF OF CLAIM 2.2. By symmetry, we can consider only the pair (x, y') . We have

$x =$ “ σ_1 would not attend”,
“ σ_3 would only attend if σ_1 or σ_2 does”, ...
 $y' =$ “ σ'_2 would attend only if σ'_1 does”,
“ σ'_4 would attend only if σ'_1 or σ'_2 or σ'_3 does”, ...

If $\sigma'_1 \neq \sigma_1$, then removing person σ_1 from consideration according to the recursive algorithm in the upper bound does not make any other \mathcal{S}_i empty. This is because σ'_2 could attend if σ'_1 does, and everyone else has at least two people that could allow them to attend. Hence $\sigma'_1 \neq \sigma_1$ implies $S^* = [n] \setminus \{\sigma_1\}$, so we can assume $\sigma'_1 = \sigma_1$, which means that σ'_1 and σ'_2 would not attend. If $\sigma'_2 \neq \sigma_2$, then removing σ'_1 and σ'_2 from consideration does not make any other \mathcal{S}_i empty. This is because σ_3 could attend if σ_2 does, and everyone else has at least three people that could allow them to attend. Hence $\sigma'_2 \neq \sigma_2$ implies $S^* = [n] \setminus \{\sigma'_1, \sigma'_2\}$, so we can assume $\sigma'_2 = \sigma_2$. Continuing inductively, at each step we have either $S^* \neq \emptyset$ or $\sigma'_i = \sigma_i$. We conclude that either $S^* \neq \emptyset$ or $\sigma' = \sigma$. \square

By Claim 2.2, the set of $m = \left(\left(\frac{n}{2}\right)!\right)^2$ parity-alternating permutations induce m input pairs $(x_1, y_1), \dots, (x_m, y_m)$ all with the same answer $S^* = \emptyset$ but such that all inputs (x_i, y_j) for $i \neq j$ have answers $S^* \neq \emptyset$. If the number of bits of communication is less than $\log_2 m = \Omega(n \log n)$, then by the Pigeonhole Principle two parity-inducing permutations generate the same communication transcript, which is impossible by Claim 2.1. We conclude that the communication complexity of the birthday problem is $\Omega(n \log n)$.

REFERENCES

- KUSHILEVITZ, E. AND NISAN, N. 1996. *Communication Complexity*. Cambridge University Press.
MOULIN, H. 1999. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare* 16, 2, 279–320.

MOULIN, H. AND SHENKER, S. 2001. Strategyproof sharing of submodular costs: Budget balance versus efficiency. *Economic Theory* 18, 3, 511–533.

Puzzle: Does Occasional Simulation Enable Cooperation?

(Puzzle in honor of Joe Halpern’s 70th birthday)

VINCENT CONITZER¹

Carnegie Mellon University and University of Oxford

Please send solutions to the author by e-mail, with the title of this puzzle in the subject header. By agreement with the editors, the best solution will be published in the next issue of SIGecom Exchanges, provided that that solution is of sufficiently high quality. Quality is judged by the author, taking into account at least soundness, completeness, and clarity of exposition. (Incidentally, there is another birthday puzzle for which we still need a solution [1]!)

This is a puzzle in honor of Joe Halpern’s 70th birthday and the June 2023 workshop (“Halpernfest”) associated with it. This workshop was held at Cornell University.

Consider the following 2-player game:

2,3	0,4
1,0	1,1

By iterated strict dominance, its only solution is (Bottom, Right) with utilities (1, 1).

However, now consider the following twist to the game: There is a 50% probability that Player 1 (the row player) first gets to *simulate* Player 2, once, before making her own move. Simulating once means that if Player 2 plays the mixed action p Left + $(1 - p)$ Right, then Player 1 will observe one draw from that distribution, and base her decision whether to play Top or Bottom on whether that draw is Left or Right. Player 2’s actual action in the game is determined by a fresh draw. Player 2 does not observe whether Player 1 gets to simulate him or not (but knows that this happens with 50% probability).

Conceptually, there are different ways to think about the game. One is the following: Player 2 needs to choose some value p beforehand, once. This value (not observed by Player 1) is then used to draw Player 2’s simulated play (if there is simulated play), and again (i.i.d.) for Player 2’s actual play.

(a1) Is there an equilibrium in which utilities (2, 3) are always obtained? Explain. (It helps to think of this as an extensive-form game.)

Another way to think about this is that Player 2 does not choose a value p beforehand, but rather at one or two points “wakes up” and has to decide which mixed

¹Thanks to Vojta Kovarik, Caspar Oosterheld, and Emanuel Tewolde for helpful feedback.

action to play (i.e., choose a p), not knowing whether he is in a simulation or not. (But, if he is in a simulation, then he still wants to play in a way that maximizes his expected utility in the real world. Also, if Player 2 plays in the real world after having been simulated, then of course he does not remember being simulated. That is, Player 2 cannot distinguish any of his different kinds of awakening from each other.)

(a2) Waking up as Player 2, what probability would you assign to being in a simulation, and how would you play? Does your answer to the above question change?

(b) Now answer these questions again for the following game:

2,3	0,5
1,0	1,2

REFERENCES

Vincent Conitzer. Puzzle: The AI Circus. (Puzzle in honor of Tuomas Sandholm's 50th birthday.) *SIGecom Exchanges*, 17(2):76–77, October 2019.